NMA 6

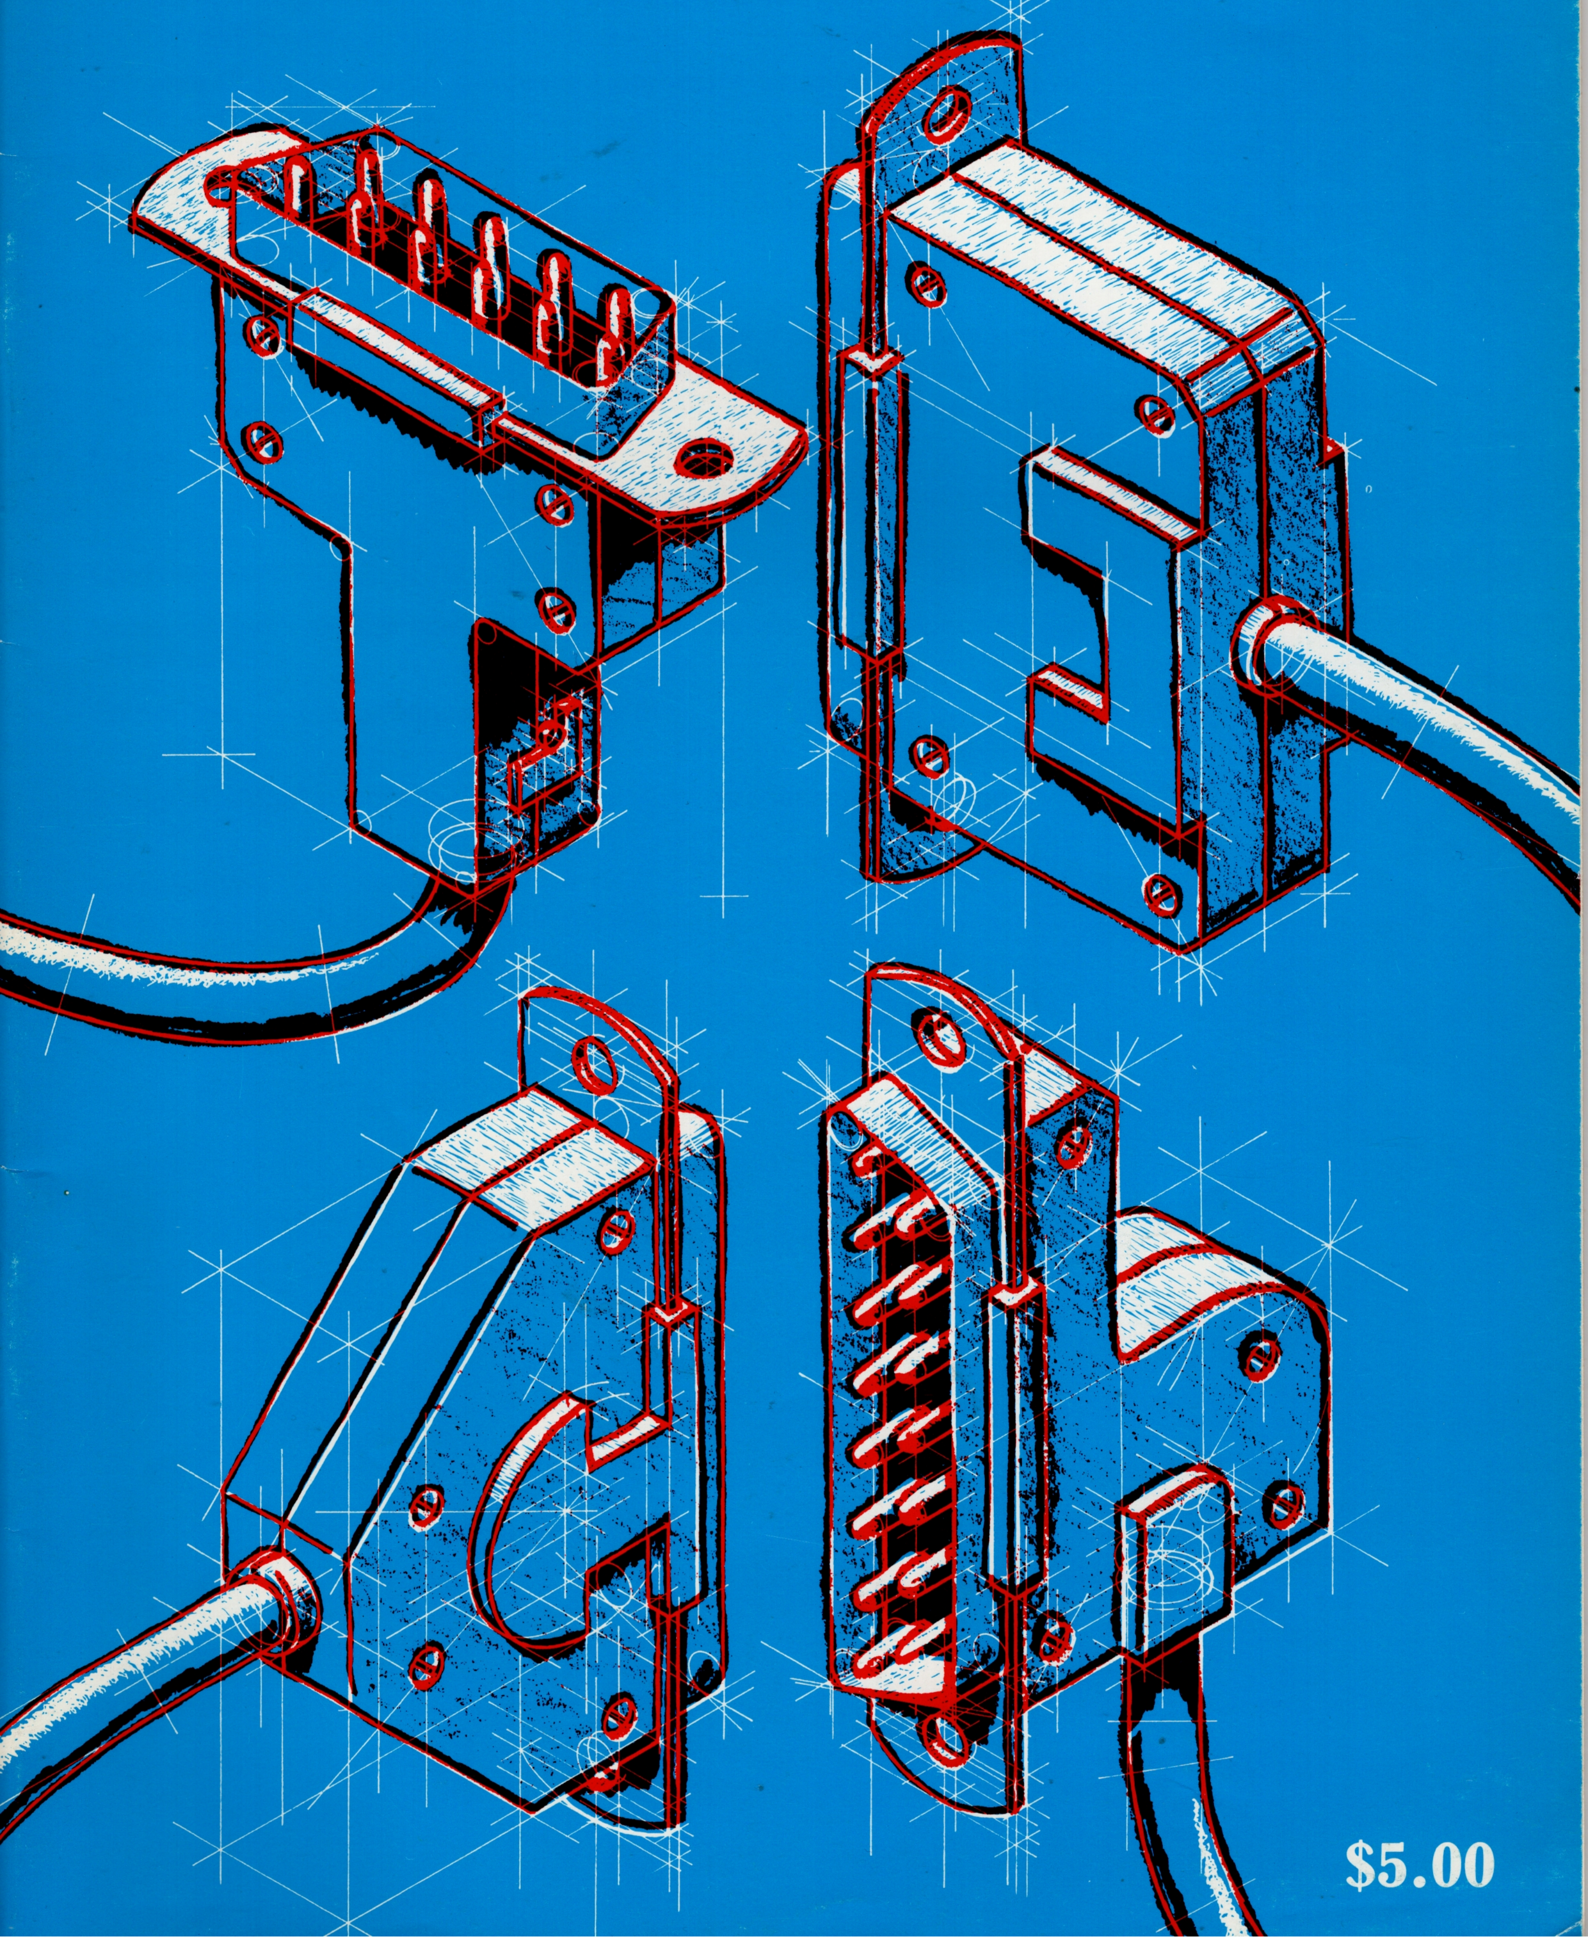

$5.00

Contents

Computer music aesthetics: a composer's predicament
Amanda Baker 3

Samples III for computer processed orchestra sounds
Warren Burt 7

Sounding & sampling
Philip Brophy 15

Towards a virtual piano action
Alistair Riddell 18

On not having read *Think — or perish! Towards a confident and productive Australia*, by Donald Horne, Commission for the Future Occasional Paper
Chris Mann 24

Beyond the body
Stelarc 27

Monophonic Variations
Greg Schiemer 31

Notes on sociological and methodological stance with respect to music
Mark Rudolph 37

Computers — performance of, and with
Cindy John 39

A digital signal processor for the real-time processing of sound
C.H. Dick 41

Perspectives
Graeme Gerrard, R.P. Harris and David Hirst 48

Editorial

The sixth issue of NMA looks at music technology in Australia through the thoughts and activities of a number of prominent composers, inventors and theorists. Although it by no means encompasses all of those active within the field, it nevertheless reflects the essential diversity and individuality present within this country.

A decade ago the term 'Music Technology' was defined by a smaller set of technical objects and activities. Beginning with analogue synthesisers which heralded a change in the sonic status quo, music technology burgeoned with the evolution of digital techniques. It brought promises and manifest changes which, while often less dramatic and immediate, have nevertheless profoundly influenced and continue to influence the course of music in the later 20th century.

During the late 1970's the potential of computer music was being realised in several prominent music departments on the east coast. In that nascent period computer music's focus was primarily the synthesis of sound. The resources available for this task were nowhere near as specialised as they are today and composition usually depended on the institution's large centralised computer systems. This was a challenging time and in their articles, Graeme Gerrard, Rex Harris and David Hirst discuss the context, characters and legacy of that period.

Over the years, institutions have continued to foster interest and provide a pedagogical foundation through which music students can come to terms with the fundamentals of music technology and its overall relationship with existing musics. The new diploma course at La Trobe University, for example, is a significant statement about the status of music technology in Australia today.

Amanda Baker and Cindy John, active with both music technology and traditional compositional practices, provide an insight into their recent work.

The domain of music technology is not, however, exclusively characterised by the activities within University music departments. Neither should it be thought of as pertaining to a single musical style, such as that heavily promoted by the commercial music world. Although these milieux have considerable impact on both thought and action in today's contemporary music they by no means define the only creative environments possibility for an encounter with music technology.

During the 1980's, the social impact of digital technology was felt across breadth of society. As its pervasiveness grew in daily life, it simultaneously expanded beyond the simple notion of 'computer music'. Innovation and exploration of new technology has always attracted individuals, many of whom have personally funded and built their own systems. This side of music technology is represented in the articles by Greg Schiemer, myself and Warren Burt.

Warren Burt who has for many years been associated with various levels of technology discusses a recent work based on digital samples of the Adelaide Symphony Orchestra performing his musical excerpts.

Philip Brophy's article presents a personalised view of music technology and the recurring theme of man/machine interfaces, while Chris Mann contributes a non-review in his inimitable style.

It is seldom that anyone outside a commercial enterprise has the opportunity to design and develop sophisticated digital hardware. Chris Dick's article provides an insight into such a project. As digital recording becomes commonplace the roll of Digital Signal Processing will take on far greater importance, allowing real-time transformation of signals. Mark Rudolph's article is a personal view of the possibilities inherent in DSP techniques.

The only article to come from outside Australia is by the artist Stelarc who now resides in Japan. Stelarc was recently in Australia and in his article he discusses his activities and current performance practices.

Finally, it is disappointing that the Sydney/Melbourne axis still dominates the representation of authors in this issue, despite attempts to encourage those in other states to contribute.

Accompanying this issue of NMA is a cassette of music from the contributors. The music is representative of the diverse possibilities of contemporary music technology and is an indication of future directions.

Editor Alistair Riddell
Typesetting BP Typesetting
Printed at Backyard Press
Published by NMA Publications
PO Box 185
Brunswick 3056 Vic
Cover Design and layout Ernie Althoff
Cassette Mastering Graeme Gerrard
Special Thanks David Rae, Rainer Linz, David Burns, Graeme Gerrard

ISSN 0811-7497

Subscription Rates
Magazine only $10.00 (2 issues)
Magazine & Tape $25.00 (2 issues)

The publication of this magazine has been assisted by the Victorian Ministry for the Arts.

Computer music aesthetics: a composer's predicament

Amanda Baker

Virginia Woolf wrote in *A Room of One's Own*:

> The woman composer stands where the actress stood in the time of Shakespeare. Nick Greene, I thought, remembering the story I had made about Shakespeare's sister, said that a woman acting put him in mind of a dog dancing. Johnson repeated the phrase two hundred years later of women preaching. And here, I said, opening a book about music, we have the very words used again in this year of grace, 1928, of women who try to write music. 'Of Mll Germaine Tailleferne one can only repeat Dr Johnson's dictum concerning a woman preacher, transposed into terms of music. "Sir, a woman's composing is like a dog's walking on his hind legs. It is not done well, but you are surprised to find it done at all."'[1] (She cites Cecil Gray's *A Survey of Contemporary Music* p 246.)

Sixty years later the prejudice is not so apparent, but its insidious persistence in habits of thought, behaviour and social structure undoubtedly continue to discourage women, and make it less likely that they will be able to achieve the 'incandescent, unimpeded' state of mind that fosters artistic vision. But in the face of this imbalance, Woolf's conclusion is that 'it is fatal for anyone who writes to think of their sex.'One shouldn't try to redress the imbalance by adopting a particularly 'female' stance in one's work. This is really the view I take. I don't write, live or work from a self-consciously 'female', 'Australian', '1980's', 'Melbournian' or 'avant-garde' point of view. I don't want to write about 'Women and composing with computers' or 'Computer music in Australia'. The battle against discrimination goes on, the struggle for national identity likewise, but I don't want to use my musical work as the battleground. Hence the following is simply a description of my thoughts about computer music in general, which inevitably have been influenced by my time, place and sex.

Frustration and disillusionment seem to be common feelings among people who involve themselves in computer music, and we spend a lot of time blaming our isolation, our institutions, our resources etc. We tend to feel that given the resources of IRCAM or CCRMA we'd be much more productive, which is probably true. Most places in Australia are still small, with fairly limited equipment and teaching facilities. I'm sure this doesn't help. We will tend to be reinventing the wheel a lot of the time, and our progress will tend to be slow for want of adequate feedback from experienced people.

To some extent we could improve matters by giving each other more assistance and feedback than we generally tend to do, but it seems that we somehow lack the courage of our convictions, and often feel trapped by gulfs of difference between our own outlook and everyone else's. Perhaps we are really more trapped by the uncertainties of the medium itself: confusion about what it is and why we want to use it. Certainly this sort of field attracts a diverse assortment of people with very different reasons for wanting to get into it.

On top of the isolation, there are big technical difficulties for musicians to come to grips with: selling our souls to MIDI with the current generation of synthesisers doesn't seem to me to be a satisfactory alternative to the difficulties presented by the multidisciplinary nature of synthesis, though things are changing rapidly and music-friendly systems are getting ever closer to being within the technical and financial reach of the average composer. Meanwhile we still have to try to arm ourselves with a useful selection from the endless labyrinth of programming skills, psychoacoustic knowledge, acoustics, synthesis techniques, etc. in the hope of being able to make sensible and effective use of the machines we have now.

In this environment of rapid change we are frustrated by perpetually being able to foresee possibilities that haven't quite arrived. Our visions tantalise us always from the next ridge, beyond which we expect to find their realisation spread out before us: the moment we lay hands on a tape recorder we think of what we could do with two tape recorders, and the moment we have those we think of what we could do with digital recording and editing. The moment we have flexible computer control we want real-time manipulation and live-interactive performance capabilities. (It is especially frustrating knowing that these things are easily being done at a few places like IRCAM, with their greater resources of machines and people and faith.) But I think the reverse of this syndrome is true, too: if beauty is in the eye of the beholder, strength is in the hand of the user; and though each new advance seems to be about to provide the key we've been waiting for, the obstacle is actually our

own inability to incorporate what is currently possible into our musical thinking. Technology will keep offering bigger and better things, while we still haven't understood some of the ramifications of the simplest first steps. In this situation survival depends on visions which can respond to the real resources we do have without being enslaved or frustrated by them.

This may all seem rather pessimistic. I think we are still suffering here from being an isolated minority within the field of current music which itself is still something of a fringe activity. But on the positive side: a recent talk by Jonathan Harvey in Melbourne brought a welcome breath of fresh air and optimism which we seem to have been lacking. He gave the impression that computers really are becoming widely used and usable by composers in Britain, and that their use is influencing general musical currency and becoming part of the mainstream. I find it very encouraging that he represents the possibility of being expert in both composing with computers and composing for conventional forces. This possibility is increasing with better machines and better programs and a growing body of music as example to learn from and follow, so that the aspiring computer musician doesn't have to set out to discover it all for themselves in quite such uncharted territory as it has been. But there are still plenty of problems.

Beyond the isolation and technical difficulties are the artistic appeal which attracts people and the artistic problems which frustrate them, as opposite sides of the same coin. Each for our own reasons we are attracted to a medium which throws open the defining constraints of our tradition of music and brings a set of bewildering new possibilities. For example: we may be attracted by the possibility of being — as a painter or sculptor is - the sole creator, with the ability to assemble the work in its final form by a process which uses direct aural feedback along the way. This situation has ramifications. It changes the listener's perspective, it removes the 'action' of performance, and the role of time in performed music, etc. etc. After trying it, some people conclude that the effects of this shift are disadvantages, which may lead them back to performed music or mixed live-electronic music. But for others those effects are precisely what they were looking for, essential to their aesthetic. The musical use of space is an aspect which seemed to have great potential. Some people feel now that that was a misjudgement:

> Much used to be made of the control over apparent spacial location which computer music provides, especially using four loudspeakers. Early pieces by John Chowning and others demonstrated that under carefully controlled conditions an effective illusion of a moving sound source could be produced... Even though I was compelled by it when I first heard it, I am no longer convinced that this aspect of 'moving sound' holds as great a promise of added expression as other unique aspects of computer music even when the effect is perfect.[3]

I think the venture into electroacoustic music involves very big and quite sudden shifts which we are partly conscious of and consciously directing, but which we cannot be fully conversant with at this early stage. Each person pursues their own path directed by an aesthetic position which they attempt to meet with a mixture of conscious understanding, intuitive nose-following and learnt techniques. It is often difficult to realise the 'promise of added expression'. The new attractions seem very exciting: for example simulation of natural sounds and the possibility of smooth transformations between sounds with quite different identities, spacial effects, juxtaposition of familiar and alien, and spectral exploration and manipulation. Such things open up endless reservoirs of new expressive effects. In trying to put them to musical use we need to try to understand what the shifts entail, though it can be argued that it is the historian's job to describe what's actually happening while the composer forges ahead guided by a blazing aesthetic vision. But surely it is helpful for composers to understand consciously what their intuitive sensitivities are telling them.

Taking up the offer of synthesis we move into the realm of playing with sound quality, focusing directly on the nature of sound and/or its associated image if it is familiar sound. Using sound for its direct absolute nature is very different from using its physical relational properties within a communally-developed system like tonality. Now it is directly laden with meaning instead of depending on its context within the 'system' (which can be very powerful with its foundation of unanimously understood vocabulary and syntax). But this difference fits with a general trend in music from the heyday of tonality and its use of sound as pattern, towards increasing focus on human gesture, demonstrativeness, self-representative expression and the sound as extension of the self; finding an ultimate extreme with voice simulation where the self-representation is absolute, the image is direct. But now it is outward looking in that it is outside the person, and exterior image of the person replacing the emanation from the person of their interior nature, or 'self-expression'. This fits with the outwardly inclined aesthetic of people like Machover and Harvey, looking for universal unification or ego-less transcendence. And while such directly self-referential sounds can draw the listener into an intense identification with them, they are not egotistically expression-laden, the expression coming through what the composer does with them. It is what the composer can do with them that is so fantastically opened out by computers.

The manipulation, exploration and deployment of material can happen with great precision and flexibility. And it is this that answers the problem of how to make music with this soundworld, with its shift of effect and affect. While it can happen that we find the promises dazzling and exciting but fail to develop them into any sort of musical language — producing very 'effects'-driven music bound by semantic stasis — often the mode of producing and manipulating sound-material itself provides the sense of system and language.

Harvey writes:

> it is in the nature of electro-acoustic music to explore the internal acoustic structure of 'notes' rather than use notes as innocent uninvolved data in musical argument. And this points towards stillness, or movement within a unity ... 'Outer' colours are not used

to colour an 'inner' abstract pattern, rather colours and pattern are one and the same. There is no 'outer' and 'inner'. The pattern lies in the colour itself.[4]

And he sees this change as moving towards a more 'spiritual art'. His aesthetic finds the right expression in particular characteristics of the medium. And his aesthetic leads him to strong ways of producing and using sound. He describes in *The Mirror of Ambiguity*[5] the manipulations he used in *Mortuos Plango, Vivos Voco*: spectral manipulation, glissandi of partials to modulate between 'bell-tonics', rhythmically-patterned readings of the digitized bell sound segments, transition between real bell or voice sounds and synthetic equivalents, and the use of the spectrum and these types of transformation and modulation to generate the structure of the piece. The contrasts between voice and bell sounds embody the juxtaposition between living and dead expressed in the text. And he also describes 'the central image of the piece, the progression from outwardness to inwardness. .. The spectacular brilliance of the attack [of the bell] gradually transforms to the prolonged calm of the deep hum note, the last to decay.'[6] His ways of using sounds are inspired by these ideas and images which are directly concerned with life and his spirituality. And he holds that the vision or idea should lead the way, answering questions of what to do and how to do it.

But for people with inexperienced and less supple and responsive imagination, and more scant resources, it can be very difficult to match our visions with our techniques. Yet the effort to do so is a good way of exercising the suppleness and responsiveness of one's imagination. I've been attempting to do this in my work with computers, and various constraints have meant that my ideas have had to adapt quite a lot to fit the limitations of the available resources. As with many people, my initial enthusiasm for and curiosity in electroacoustic music resulted from time in the analogue studio (at Melbourne University), which I found liberating - because it afforded freedom from instrumental line and 'note', free rhythm, sound imagery, and the ability to construct by aural trial and error.

The move to computers hasn't really brought the same sense of freedom until quite recently. At first it just seemed to be an endless amount of numerical specification in terms of acoustic properties which I didn't fully understand. That was using a version of Music V; and learning the ropes of synthesis at a low level without a whole team of people to compare notes with and swap discoveries with ended up discouraging me, though the potential and the realities which other people have made of it are still very exciting to me. Then the department succumbed to the MIDI-Yamaha epidemic and I hesitantly decided to try that instead. So out went all my visions of sound-image exploration, and transformations. I've always wanted to try things like transformation between the human but vast sound of cheering football crowds (as heard from quite a way away from the stadium) and the inanimate but lively sound of breaking waves. The combination of their similar rhythms and the changing images fascinates me. But these sorts of things were beyond me.

Turning to MIDI and the TX816 I had to use largely predefined sounds and 'notes'. Since the best sounds on these synthesisers seem to be the simple ones, and those aren't very interesting, I decided to explore the possibilities of very dense streams of simple sounds, building granular textures and having precise control of their overall envelopes. And since this medium is pitch oriented to suit conventional musical use of pitch, I decided to use it partly as an exercise in controlling the 'harmonic' complexion of such textures, to see what I could use of my conventional understanding of pitch and pitch relationships. Harmony (in the general sense: resulting from different types of pitch combination) is still something I view as very important, and difficult to use well in modern music. Would this idiom — where voice-leading and the difference between simultaneity and succession don't really exist — bring a different but interesting manifestation of the effects of pitch relationships? I set up my method of generating notes so that I could manipulate harmonic colour and harmonic motion by controlling the type and complexity of pitch aggregates, the type of change between different pitch collections (i.e. smooth transition, sharp contrast, registral displacement etc.), speed of change etc. That exploration produced one piece (*Colloquy*) in which I used cascade gestures which move between randomly chosen pitch content and specifically weighted pitch content — the chaos-order idea — and it followed a global progression of growing and then dissolving contrasts between the specific pitch collections used, so that after reaching a high point of tension and juxtaposition, things were gathered up into one harmony and finally that dissolved into subsiding indistinctness. I was also exploring wave-like aperiodic rhythms, which fitted my interest in using patterns and processes like those of the natural world, which move with a plasticity not afforded by the constraints of our traditional rhythmic systems, and which we seem to hear in a less pattern-predictive/comparative way.

Reflecting on all this I have come to think that the biggest musical potential is in the contrast or juxtaposition of different things, attitudes, motions etc: that these 'outside' sorts of rhythms might provide a context for very human gestural rhythms, or that their soporific quasi-regular rise and fall could offset sporadic explosions of human activity. The same applies for harmonic things where performing forces require and suit a particular bandwidth of density of notes and hence pitches, or pacing of harmonic motion. Instruments and voices can provide nice sinuous lines of monody or spare textures where every single note's pitch and relation to its neighbours can be important to the motion. In the computer textures the significance of single notes and single pitches is reducible as the density increases, to give very subtle shadings. One can play on the continuum between many and therefore individually weak contrasts and few and therefore strong contrasts, sliding between very complex subtly shaded harmonic collections or motions and very elemental material where each single note/pitch etc. is strong and structurally influential.

Relating contrasted elements, either by juxtaposition or transition/transformation answers an aesthetic desire which I think many people have, of relating, tying together, making sense of our

fragmented view of the disparate world we live in. That is something Jonathan Harvey talks of, saying the interest is more in relating widely different things than producing a tightly unified single minded sort of music, and he is fascinated by the process of transition from one identity to another. Tod Machover also writes about this in *A Stubborn Search for Artistic Unity*.[7] I've been hoping to relate the outward feel of wave-rhythms, super-human cascades of mingled sounds, aperiodic rhythm patterns etc., whose precise subtleties we pick up and respond to even though we can't 'play' them easily, with human gesture, familiar sounds — directly personal expression — as a way of presenting an outward but human sense of contemplation. My use of pitch is part of this attempt. I think harmonic nature is something we latch onto and identify with in a strong direct way, and can become the 'subjective' agent in very unhuman material. A cascade of lovely sounds can be effective and attractive. Giving it a calculated harmonic shape can bring it to life, and reiterating it by shifting or developing its harmonic nature seems to help to bring a particularly human expressiveness.

Another piece I've been working on has introductory and concluding sections for tape alone, framing four movements for instrumental ensemble and choir (*Birdsongs*). I've tried to use a sort of complementation so that the different media offset each other without seeming dislocated or unrelated. The texts give a humanly fanciful, fairly anthropomorphic contemplation of particular birds, and through them the world in general. The tape part uses the unhuman sounds of synthesisers and the animate but not human sounds of birdsong. The live part becomes the human response to these. It is all human in that it is all supposed to be music, but the tape parts are more passive, and somehow beyond, in that the material is mostly completely different from humanly performed music.

The position of birds and birdsong is peculiar. It is 'outside' rather than human but we have imbued it with all sorts of images and feelings so it has become symbolic. Having evolved as a means of communication, birdsong has a coherence and consistence, and being aural and pitched it can be perceived 'musically' because of our predisposition to make our own sense of sounds that flow and are pitched or have some other strong characteristic. So we can have our own genuinely musical and meaningful reception of those sounds, at the same time as finding them bewildering at some level. The 'real' meaning of the message (alarm, territorial, mating etc.) is inaccessible to us: it is not our language, and that strange sense of distance can be part of the meaning we impose on it. So there is a triple perception going on:

- the extra-musical associations birdsong evokes for us (the country, Spring, playfulness, whatever!, and our anthropomorphic ideas about the birds and their characters),

— the strange sensation of detecting its coherence yet not understanding it in its own language,

— and our perception of it because of that coherence in our own musical terms.

My piece tries to use birdsounds, computer-generated material and choral/instrumental material in a complementary way that will tie these different points of view: using the natural magpie call (familiar yet alien), the computer soundworld (alien yet familiar in some of its gesture and harmony), and the live performers. I have tried to tie them by reflecting the birdsounds, the birdworld, the world, and human feelings about them in all three media. Its not simply that I want to represent these differences, set them up or point them out. I can do that in words. It is that I want to inhabit their relations of contrast musically. I always find I need to develop a sense of the 'nature' of the ideas and elements I am using. That nature can end up manifesting itself in all sorts of ways, which produce the particular surface characteristics of the music. That's why I find it difficult to plan a piece in terms of "slow chord section here, using specific pitch deployment mechanism X here, with rapid sporadic interruptions here and a climactic fusion of motives B and E here..." These things only emerge from a process of exploring the elements, ideas, and sorts of contrasts, definitions, relationships they generate.

For my instrumental sections I ended up defining quasi-magpie gestures using particular pitch material that was **not** like the actual pitches of the magpie calls I used for the tape part. In the instrumental context it gave a better impression of the humanly perceived 'essence' of the birdcall than a clarinet playing the simple pitches actually found in the birdcall, and it provided more interesting musically workable pitch material, so that the gesture became a musical element, not just an instrumental 'quote' of the bird. The simple pitches of the magpie song were used when the real recorded magpie sound was used, in the tape parts. Here the call was artificially manipulable (I had it on a sampler) so its simplicity was a good starting point. From there it could be transposed and fragmented. I used it to define a harmonic anchor for the computer generated sounds that it was combined with, and transposing it initiated warp or departure/intensification. The mellower unreal sound of the slightly lower slowed birdcall seemed to fit this harmonic stretching. Its simplicity also seemed to fit well with the subtly changing harmonies of dense computer generated textures. So: (I hope) the oppositions between the situation of the real birdcall, the human perception of it and the human instrumental 'birdcall'-response become musical, and the operation of musical material grows from this complementary rather than similar reflection between tape and live sections.

The sorts of contrast, distinctions, dichotomies, juxtapositions, ambiguities, combinations that we use, we use for their nature, that can become the nature of the music and somehow represents the nature of the vision. We don't only use a juxtaposition or pattern because it is pleasing, but also because it enables the material to take on meaning and coherence. And if that juxtaposition or pattern or relationship does become the musical nature, then its original trigger (the model or image in the composer's mind) isn't essential to the awareness of the listener.

Ultimately I can't tell if my piece 'works'. Maybe the images of the birdcall don't really fit semantically into the musical flow. This is the sort of question that I suppose is being gradually answered as explorations fill in a clearer picture of the 'musical'

properties of material and methods now available. And this is what makes it all so fascinating and so frequently frustrating.

Computer music may still not be central to twentieth century music, but it takes one invariably to the heart of many issues which are: time, timbre, harmony, the role of performance, the nature of gesture etc. So it is relevant and rewarding in general even if it is still an eccentric way of approaching those matters. The unique angle of its view of these issues suits a peculiar sort of aesthetic, which may always be a bit alien to the mainstream, but which appeals very strongly to some composers. And some people do see it as more and more part of the mainstream: Machover concludes his chapter by writing

> In the more artificial world of music we are perhaps closer ... to being able to manipulate both computer and acoustic materials in a single musical context, with no difference between them, the one enriching the other, totally compatible, totally unified: one single artistic vision motivated by deep human concerns.[8]

From that point of view both elecroacoustic and instrumental media are integral to a modern musical outlook.

Footnotes

1 Gray, Cecil. **A Survey of Contemporary Music** London: Penguin, 1928. p 56
2 Ibid. p 102
3 McNabb, Michael. *Computer Music: Some Aesthetic Considerations.* in **The Language of Electroacoustic Music** ed. Simon Emmerson, Basingstoke and London: Macmillan 1986, p 147-148.
4 Harvey, Jonathan. *New Directions: A Manifesto*, **Soundings** Winter 1983-84, p 10.
5 **The Language of Electroacoustic Music**, p 181ff
6 Ibid., p 181
7 Ibid., p 191ff
8 Op. cit. p 215

Samples III for computer processed orchestra sounds

What it is and what it is not

Warren Burt

I want to tell the story of how a piece of mine came to be written,mixing up technical details, aesthetics and anecdote much in the way they occur in life. To explain both the how and the why, the practicalities and the philosophies, roughly in the order that they happened, or that memory allows me to think they did. To show that for this composer, at least, working with technology and grappling with aesthetic concepts go hand in hand, indeed, are inseparable.

Since earliest childhood, I have been thrilled with the sounds of the recorded orchestra. For me, the primary reality of orchestral sounds has been that of the recording. Listening to a recording of an orchestral performance has been for me in no sense a simulacrum, but a strong and resonant reality with its own characteristics quite different from those of a concert hall performance. Both sonically and sociologically, I have come to prefer the intimate experience of hearing the orchestral recording to that of hearing 100 musicians performing in front of several thousand people. Indeed, during the recent visit of Olivier Messiaen to Australia, I stayed home and heard the broadcast of the Melbourne Symphony's performance of his *Turangalila* in my living room. I do not feel I "lost out" on anything by choosing to experience the work in this way, in fact, enjoyed it more as a radio piece than if I had been in one of my least favourite places, the Melbourne Concert Hall.

I have always wanted to write for orchestral sounds, and have done so three times in the past. *Drakula* (1969) for large orchestra was performed at a reading by the Buffalo Philharmonic Orchestra, and the more experimental *Face* (1973) was performed by the La Jolla Chamber Orchestra. However, with my moving to Australia in 1975, all chance to have access to the orchestra disappeared. Unless I was willing to write within the

limitations and expectations of orchestras here, it was quite clear that that resource would be closed to me. Further, since most of the work I was doing involved the investigation of new musical sociologies and performing situations, it was clear that those interests were incompatible with working with that most sociologically fossilised of musical organisations, the orchestra. So I was quite content working with other media.

I started working with sampling in 1980, on the Fairlight CMI, and quickly realised that many of the tape manipulation techniques we had been doing with cassette recorders, in both solo and Plastic Platypus group performances, were immediately possible with sampling. In 1982, I wrote *Dedication Canons* for string orchestra, and, having absolutely no access to an orchestra, realised it on the Fairlight. Sampling techniques then were not what they are today, however, and I remain less than satisfied with my realisation of this piece. It was not until I had a sampler of my own, in 1986, that I began toying with the idea of making polyphonic textures out of samples longer than a single note. With my first sampler, the toy Casio SK-1, one could have a single sample 1.4 seconds long in glorious lo-fi four voice polyphony. Just for a lark, I sampled a cassette of the Ravel *Piano Trio* that was handy. The first sample was of three ascending thirds on the piano. I quickly played the keyboard of the sampler, hearing, at various speeds, the recording of those three dyads each time I pressed a key. I pressed many keys quickly. Voila! Instant texture! Polyrhythms, polyharmonies and thematic unities galore! More orchestral material was sampled. Oboe and string textures from Delius. Flute and string ensembles from Glazounov. Piano textures from Debussy. All of these were recorded on a four track recorder and mixed. The resulting sound world was pretty close to the kind of orchestral writing, filled with simultaneities and juxtapositions of different textures, that I had always wanted, but which was clearly impossible in the Australian situation. I was delighted. While nothing I had done would have been impossible using classical tape techniques, the ease of doing it with the sampler made it much more immediate. This short piece, *Samples for Orchestra*, became the first movement of a longer piece, *Lo-Fi Proposals*, which explored a number of the possibilities of the Casio SK-1, and was performed at the Experimental Intermedia Foundation in New York on Oct. 21, 1986. However, I gave *Samples for Orchestra* the subtitle, *"neither a deconstruction nor an appropriation, neither bricollage nor post-modern,"* because I wanted to make it clear that in this case I was not playing any of the games with found material that were part of much contemporary art practice. I have played these games many times since the late 60s, and continue to enjoy playing them, but here, I was using found orchestral material only because I had no other access to the orchestra. Right from the outset of this project, my aim was to abstract the material into interesting polyrhythmic, polytonal, and polytimbral textures that would be impossible for live performers, avoiding, as much as I could, the sense of the material being pirated or quoted.

The issue of pirating, or quotation, that is implied by the very nature of the sampler would not go away, however. And so, when I acquired my AKAI S900 sampler in late 1986, one of the first things I set out to do was to make a live performance piece made exclusively of quotations from recordings of orchestral works. In contrast to the quotational promiscuity of *Samples*, I decided to use only one composer, Maurice Ravel, as my source, and to have the quotes sometimes very obvious, and sometimes quite abstracted. The piece was a live keyboard improvisation, using six different multisampled keyboard patches such that each keyboard patch was divided into five or six discrete areas with different samples on them. On any individual keyboard, I could play a variety of quotes at different transpositions and tempi, and by switching from keyboard patch to keyboard patch could have access to different combinations of quotes at different pitch levels. To continue the controversy over the nature of what I was doing, and knowing I would be performing this piece mainly in the visual arts context, I called this piece *Samples II for Orchestra: Ravel Hommage (that which is neither a deconstruction nor an appropriation, neither bricollage nor post-modern)*. In various talks before performing the piece, I explained that though I was quoting, I didn't feel I was deconstructing Ravel, because I was not pulling any work of his apart in order to make statements about it, and that though I was indeed violating all sorts of copyrights by using samples of recordings in this way, I didn't feel like I was appropriating anything from Ravel (Pierre Boulez and Columbia records definitely, but not Ravel), because, after long study and familiarity with his work, I had come to regard him almost as "family," if you will, and felt about using his work much as I would feel about using the work of any living colleague who had given me permission, or even invited me to use their material. Furthermore, mine was a very studied use of quotation, I did not regard myself as the innocent bricolleur, assembling new works out of whatever came to hand. I have done that many times in the past, but this was different, the quotes were very carefully selected, and the ways in which they were put together was also carefully considered. Finally, I wanted to call into question the whole use of the term "post-modern," which has been so overused in describing our activities. I maintained that all the techniques I was using in the piece, although made more accessible by technology, were already present in the work done by Charles Ives around the turn of the century, and that if he was considered post-modern, then we would have to consider our whole century as part of the post-modern, and have to hunt around for a time when modernism actually existed.

Samples II was quite a successful piece. In it, I was able to play live the kind of orchestral textures I could only dream about before. However, I wanted very badly to be able to do this kind of piece with orchestral sounds of my own devising, and not just with quotations. This clearly seemed like an impossibility given the reality of orchestras in Australia, and so I shelved the idea for possible future reference.

Then, in early 1987, I was selected as one of the ABC's three composers-in-residence. As part of my application, I proposed that should one of the ABC orchestras become available for a day, due to

another project being cancelled, I would write a number of fragments for orchestra which could easily be sight read, have those recorded and then treat those with my sampler. I felt that the nature of the transformations the sampler could affect on the material meant that I did not need to have original material of great complexity. Here I could have the best of both worlds: I could work within the limits of the Australian orchestras, and thus get out of them what they could do best, and then use the technology to extend the sounds they made into textures that were well beyond the realm of any live performers anywhere. Thanks to the heroic efforts of ABC producer John Crawford, the Adelaide Symphony Orchestra became available for a single day's recording session on August 3rd, 1987. This was only confirmed on July 19, and thus I had two weeks to write and copy all the parts. When I pointed this out to John, he said to me, with a giggle in his voice, "You can do it. You're a hack." He was right, probably on both counts.

I immediately realised that this would probably be my only chance, for the forseeable future, to work with the orchestra, so I gave considerable thought as to the nature of the piece I wanted to write. My main (though not my only) desire as an artist is to create things for which there are no models, and then to use these things as vehicles for perceptual exploration, to find out what it is these new things are, how they work, and how we work, what we do, when presented with this new information. In this exploration, I wanted to work seriously with many of the issues that concerned me: the reality of radio and tape work as a medium in itself; the use of non-directional, constantly changing forms; the use of extreme duration; the use of many different random methods of composition, each involving interaction between various electronic music systems and physical activities; working with a body of material marked by its diversity, and not by any system of thematic unity; and the creation of material almost in real time, with no, or very little, second thoughts used in the creation of the work. This last was practically forced on me due to the extremely short time I had to compose the material, but I did not mind all that much, because I always wanted to work with, as Jack Kerouac says, "no fiction, no craft, no revising afterthoughts, the heartbreaking discipline of the veritable fire ordeal where you can't go back but have made the vow of "speak now or forever hold your tongue" and all of it innocent go-ahead confession, the discipline of making the mind the slave of the tongue with no chance to lie or re-elaborate" (*Desolation Angels*, 1960, p. 238). This project would give me, of necessity, the chance to do just that.

Because of my ABC residency, it was clear that the piece was to be a radio work, and I have long felt that radio composition is a distinct medium in itself, which does not need to dignify itself by comparisons with other mediums. As Kenneth Gaburo says about tape music in *Isit*, (and the same can apply to the closely related, but subtly different, medium of radio) "Tape compositions are as direct as one can get. There actually is nowhere else one can turn to which will do any "good", except literally to 'face' the music" (*I.S. Journal No. 2*, Los Angeles, 1986 p. 50). One of the main outputs of ABC-FM is recorded orchestral music. I decided that my tape/radio piece using orchestral sounds should play with this fact, and for most of the time, should sound orchestral. That is, very few of the modifications used should distort the sound in such a way that the sounds began to sound "electronic" or even transposed to such a degree that their essential "orchestral" timbral identity was lost. The piece should mostly sound as if it could be played by an orchestra, albeit a superhuman one.

And I wanted the piece to be LONG. I conceived of the piece as one possible extension on the transcendental orchestral tradition of composers such as Ives and Scriabin, so I wanted to make a piece that stretched my durational capacities as well as those of my listeners. To experiment with musical enlightenment through endurance, if you will. To make a piece that in memory could not be recalled as a single, or even as a handful, of "images." Further, to make it long as a positive statement that there are alternative ways of thinking about our attention spans than that 10 second to 2 and 1/2 minute time span the media seems to impose on us. I wanted to show that long time spans of serious activity are still available to us, no matter what most people, in and out of the media, want us to think.

Formally, I wanted the piece to be non-directional, both in its small scale and large scale structures. I enjoy structures which seem "to set out in no direction and arrive not knowing where, to come and go without knowing where it will stop," as a follower of Chuang-Tzu wrote about 2200 years ago (*Chuang-Tzu*, tr A.C. Graham, Unwin Paperbacks, London, 1986, p 162). I find these structures more in line with my ideas of exploratory music, of investigating what the effect is on each person of structures for which no models have previously existed; and I find these structures more useful politically, in that they avoid the setups of expectation of reward and defeat of those expectations that tonality practices. I have always found this kind of manipulation particularly distasteful, and would rather make work where each item is valued for its own sake, and not its hints at what is or is not to come.

On the small scale, this kind of non-directional form would mean that I would use many different random processes to generate much of the musical material of the piece. These would not be used to merely "get beyond my own immediate tastes," but, more deeply, to explore the kinds of music that would result from these machine and physical processes, in order to create work I would have to learn to listen to, to experience. On the medium scale, a non-directional form played over this duration mandated, for me, a diverse range of surface styles, so that changes and juxtapositions could be clearly heard as such. And on the largest level, it meant that the overall form of the piece would also be determined by quasi-random means, so that the progression and juxtaposition of the various materials and their treatments was as non-directional, and "non-intentional," as not concerned with cause-and-effect, as I could make it. Again, I wanted the piece to be exploratory, to be one of those pieces that would be done only when I "came to hear what it is I had made" (Gaburo. p. 48), to be a piece with many unfamiliar and com-

plex kinds of juxtapostions one would have to grapple with, creating a complex sound object which would richly repay many listenings, providing something for those who would like to work on this kind of experience.

I decided that composing the piece would take place in three stages:

1) Composing many small fragments for orchestra using the full resources of my studio, using many man/machine (and some non-machine) random (and a few non-random) processes to generate the scores. Personal Composer software was used on my IBM clone, and information was fed into it through a Casio CZ-101 keyboard, or else through control voltage input from my JL Cooper control-voltage-to-MIDI converter. The control voltages were generated by my battery of electronic music equipment, including Aardvarks IV, my homemade random control voltage generation system, my large Serge analogue synthesiser, and a Gentle Electric pitch and envelope follower.

2) Taking the recordings of those fragments into my home studio, making samples of them on my AKAI S900 sampler, and using MIDI input, either from the keyboard or from the computer or electronic music systems to process the samples into recordings of orchestral textures.

3) Mixing those recordings into an overall structure such that different textures were juxtaposed both on top of each other (vertically) and following each other (horizontally). Since the recorded textures would already be fairly thick (eight layers maximum if only one pass through the Akai was used, more if the texture was made with a multitrack recorder), I decided that the final mix of the piece would consist of a maximum of three recordings at any one time. I was interested to see if we could learn to hear, and follow, even this many discrete complex textures as a simultaneity.

Fifty short fragments were written, for everything from full orchestra to various chamber combinations to solos for various instruments. The maximum length of any one fragment was about 22 seconds, with the majority being much shorter than that. This upper limit was determined by the longest sample I could get on the Akai at a sampling rate that I felt still gave me enough bandwidth for FM broadcasting. Longer samples were possible, given the Akai's infinitely variable sampling rate, but at a trade-off of frequency range I did not feel like making in this piece. Eight different kinds of material were written:

1) Six fragments for full orchestra, each of which was generated by a different interactive random process. In each of these, and in all cases of polyphonic textures made in this piece, each individual line was generated separately, without referring to the other lines, the juxtapositions resulting by chance.

2) A series of wind, string and brass chords, each of which was a different voicing of a D-F-A-C-E-G chord. The wind voicings favoured perfect fourths, the string voicings, fifths, and the brass voicings thirds. These were recorded as both staccato attacks and as sustained chords, and were used to make "single-note" samples to assemble "beds" of sound and other textures from.

3) A series of random chords where each member picked their own pitch for each new attack.

4) Eight polyphonic fragments for chamber groupings within the orchestra, derived, like the full orchestra fragments, by a variety of random processes.

5) Ten single line melodies doubled at the unison or at some interval of transpostion (i.e. the same notation given to a non-transposing and a transposing instrument results in a doubling at the interval of transposition). These melodies were generated by keyboard improvisation, serial methods and in two cases, just plain writing down what I heard in my head at the time. The choice of doublings (were there two or three instruments playing, what were they, what tempo were they playing at?) was made by referring to a chart of random numbers.

6) Improvisations based on limited timbral resources, such as the fragments for trombone glissandi, or the percussion duet fragments.

7) Solo line melodies, written either by performing on Aardvarks IV, my homemade random voltage generator, or by improvising on a keyboard with my eyes shut and the sound turned off, so that the resulting music would be as pure a result of my own physical gestures, without sound feedback of any kind, as possible.

8) Single notes for various instruments, designed to be used as single-note samples in the traditional sense, to be used in assembling microtonal textures.

Additionally, each fragment was recorded at at least two different tempi. All the fragments were in 4/4, with only quarter, eighth, and sixteenth notes used. Triplets and other divisions of the beat were not felt to be necessary, given the nature of the processing. Six tempi were used, and these were related by the proportions of a justly tuned D-F-A-C-E-G chord: 60, 72, 90, 108, 135, and 162 beats per minute. This meant that a fragment recorded at say, 72, could be played a minor third lower than recorded, and it would be playing at approximately 60. When juxtaposed with the original 60 beat per minute fragment, this would result in two fragments in rough rhythmic unison doubled at the minor third. Or conversely, if a fragment recorded at 60, and the same fragment recorded at 90 were played simultaneously, a rhythmic canon at the unison would occur, where the tempi of the two fragments would be related by 3/2. In the course of the piece, most kinds of rhythmic/pitch canonic relations occur.

To give an example of the kinds of processes used in composing some of the fragments, here is the procedure I followed in composing the polyphonic chamber music fragments, material type 4:

1) Is the fragment 2, 3, or 5 voiced? Pick a random number (from *A Million Random Digits* Rand Corp. 1955) to determine

2) Which instruments are used? Pick a random number from 1 to 20 for each voice, with each instrument type in the orchestra assigned one of these numbers.

3) Will it be played at two or three tempi?

4) Which of the six tempi will each of these be?

5) Which composing method will be used? There are six possibilities here (which were also the six methods used, one each, to create the six fragments for full orchestra). Random numbers will determine which of the six are used.

Ex 1. Fragment 11: Chamber Music 1 *heard from 2'23 2'41 as a solo, from 21'57 23'09 as a unison chord, and 22'12 25'02 as a polytemporal unison canon.*

a) Free improvisation on the MIDI keyboard, which feeds into the notation program. In this process, as in all the others, each individual line is made independently of the others (previous lines are not listened to), with the resulting juxtapositions left to chance.

b) Aardvarks IV generating random control voltages which are fed into a control voltage to MIDI converter and thence into the notation program. The notes will have varying durations and hold times. With this, and the following two processes, my performance controlled pitch range and tempo, but not the moment to moment details of note generation.

c) The Serge analogue synthesiser's Smooth and Stepped Function generators patched to generate a chaos like stream of voltages. These result in continuous notes (i.e. very few rests) when fed through the control voltage to MIDI converter and into the notation program.

d) Aardvarks IV with all short duration notes and each voice accelerating and decelerating independently.

e) A tape of Canter's Deli, Fairfax Avenue, West Hollywood, on a particularly busy Friday evening, fed through a Gentle Electric Pitch and Envelope follower and the result fed into the control voltage to MIDI converter and the notation program. For b, c, d, and e, the control voltages were patched through the Serge Control Voltage Processors, allowing me to specify the ranges of the control voltages so that notes out of range of that particular instrument would not be generated.

f) Random numbers chosen from *A Million Random Digits* played in via the MIDI keyboard input. These could be applied to both diatonic and chromatic pitch sets.

Example 1, Chamber Fragment 1 shows an example of the type of output generated by this process. It is a 5 voiced fragment for oboe, clarinet, bass clarinet, tuba and solo violin, recorded at both 60 and 90, and was composed using the Serge patch described above in letter c. I was especially surprised at the "Schoenbergian" sound of this fragment, delighted to be reminded that when one is working with such a "stylistically saturated" medium as the orchestra, references are bound to continually occur, no matter what methods one uses in composing. This was one of several fragments generated randomly that sounded as if they were references to the "style" of another composer, and I happily accepted them as an innocently delightful reversal of the composing situation of the first *Samples* piece, where I had tried to make quotations sound as if they weren't. Here, there was no quotation whatever going on, but occasionally, things sounded like they might have been. As Ives was fond of saying, "What music sounds like may not be what it is."

The score and the parts were finished on time. Thank God for computer music notation systems. They may be funky, clumsy, limited, and have more bugs in them than a case of two year old flour, but in the end they made this piece possible.

The recording session was a total delight. The Adelaide Symphony, the conductor, James Ferguson, and the producer, John Crawford, were all wonderful. I've seldom had such polite, enthusiastic, co-operative, friendly working conditions with any group of musicians as I had on this occasion, and I am incredibly grateful for the wonderful time I had and the very good performances I got. We recorded almost all the fragments I had written (I was relying on the pressure of time to winnow out those few fragments I felt I really could live without- another form of chance technique for selection and winnowing of material- how much can we get recorded in the limited time we have?), and I left Adelaide with almost two hours of material on tape, which included several takes of each important fragment at a number of different tempi. I had discussed the process of the piece with the producer, the conductor, and the concertmaster, and they were all sympathetic to it. They knew knew how to get the accuracy as I desired out of the orchestra, and also when to stop trying to push for nuances that would not survive the processing. Strangely enough, though I prefer hearing orchestral sounds through loudspeakers, here I preferred being in the studio with the orchestra, being physically present with the players, making suggestions, etc. rather than sitting in the recording booth. It felt both friendlier and more intimate to be with the players while they worked, and also felt as if I was in more immediate contact with the sound as it was being generated. The session really was one of the most rewarding musical experiences in my life.

On my return to Melbourne, I set about experimenting with the material on the sampler. It was obvious that if I was treating one sample at a time in the sampler (a limitation of my own, not of

the sampler's) or samples of the same fragment played at different tempi, I was dealing with canons. These canons could be of a number of types, as outlined above. Other processes were also evolved. In all, five different kinds of processings were used:

1) Canons of various kinds. These were of two types : untransposed canons and transposition canons.

a) Untransposed canons were made by taking a single fragment and sampling it, usually not making a loop out of it, and then setting the keyboard so that the same sample at the same pitch appeared under eight different keys. If any of the keys was pressed, that fragment would play. If one removed one's finger from the key, the fragment would stop playing at whatever point it was in it's playing. In this kind of canon I did not want to work with partial fragments, but with their whole length, so having the fragments without loops meant that I could simply leave my finger on the key for a duration longer than that of the fragment, and the fragment would simply stop sounding when it reached it's end, even if I left my finger on the key. This was quite handy when you consider that I could have up to eight identical samples playing at any one moment, and telling which fragment was under which key could get to be quite confusing. There were several different forms I used for the untransposed canons, but one of the ones I used the most (in fact, with hindsight, I feel I may have overused it just a little) was where I took the tempo of the fragment as a guide, and brought in each of the eight fragments at a distance of 4 beats from the start of the previous fragment, then at a distance of 2 beats, and finally at a distance of 1 beat. Since the fragments were usually 8 or 16 beats long, this gave a series of canons which got closer together, and denser, as they progressed. Also because of the small amount of material being used, the canons became highly repetitive, almost fugal in nature.

Another kind of untransposed canon was made when each of two or three recordings of a fragment made at different tempi was made into a sample. These were made into loops, and all of these were present simultaneously inside the sampler. Each was assigned to a different key on the keyboard. When these were pressed simultaneously, and held down, the samples would play simultaneously. If, for example, a solo clarinet line beginning of "C" were recorded at 90 and 108, and samples of both recordings were played simultaneously, both tuned so that they, too, began on "C," one would get a two part rhythmic canon at the unison, with the two voices in rhythmic relationship of 5:6. This two voice canon would eventually come full circle, and start repeating, after 6 repeats of the faster voice, and 5 repeats of the slower one. One complete cycle of this canon would be used as a recording in the final mix. If the samples were of a single line or a doubled line, the result would be a quite simple polyphonic texture. But if the samples were of a complex orchestral or chamber texture, the results could be quite thick and dazzling.

b) Transposition canons could be of many sorts. As discussed earlier, they could involve fragments at different tempi transposed to be in rhythmic unison, but doubled at various intervals, or they could involve one or more fragments played at different pitches than they were recorded at, making polytonal and polyrhythmic canons. For example, a recording of an orchestral fragment beginning on "C" is played simultaneously on two keys a major second apart. One is tuned to play the fragment beginning on "C," and the other beginning on "D." Since the ratio of a major second is 9:8, not only will the higher sample be playing a major second above the lower, it will be at a faster tempo, related to the tempo of the lower sample by the ratio of the pitch transposition, in this case 9:8. If the two keys are held down for, in this case, 9 loops of the higher pitched sample and 8 loops of the lower pitched sample, the canon will come back to it's beginning. Thus, a large variety of polytonal and polyrhythmic canonic relationships was available through playing samples at different pitches, and this gave me not only textures not available through human abilities, but also, on a more mundane level, timbral resources not available due to economic limitations, as in the canon for 5 piccolos heard near the start of the piece. The danger here, however, was that if one played the sample transposed by too great a degree, the sample would lose its "timbral identity," and begin to sound like a speeded up or slowed down tape, and not like an orchestra. Generally, I found that I could not transpose most textures by more than a perfect fourth before they began losing their identity. In a few cases, the margin of tolerance was much greater than this (with bass clarinet, for instance), but in other cases (some strings, some brass, english horn, etc.), it was actually much less. This limited the kinds of transpositions I used, as did my desire to keep the recording resulting from these samplings to generally under three minutes each. A doubling at the minor second, for example, or at some microtonal interval, might result in a complex rhythmic relationship of 15:16 or greater, but the time taken for a complete cycle of this relationship would usually be inordinately long. Simple relationships, such as those implied in the D-F-A-C-E-G chord, generally proved to be more useful for this piece.

2) Microtonal textures obtained through detuning samples by as little as 6.25 cents. I called these textures my "Niblocks," after New York composer and film-maker Phill Niblock, who has worked for many years with gorgeous micro-tonal drones made of recordings of instrumentalists playing single notes tuned with extreme precision. In some cases, a single note or chord would be sampled and made into a texture, and in other cases, it would be a melody or texture that would be so treated. This was the only case in which I departed from my rule that the sound output should sound "orchestral." Some of the phasings and beat patterns that result from these treatments sound quite "electronic." However, I justify the use of these in two ways- first, their "close but not quite orchestral" nature does make apparent the nature of the piece, one made with orchestral recordings, by flirting with the boundary between "orchestral" and "recording processed" sound, and second, I wonder how many of these effects we might have obtained if indeed we had spent weeks training players to play their instruments with that precision. That is, I'm not really sure how much phasing,

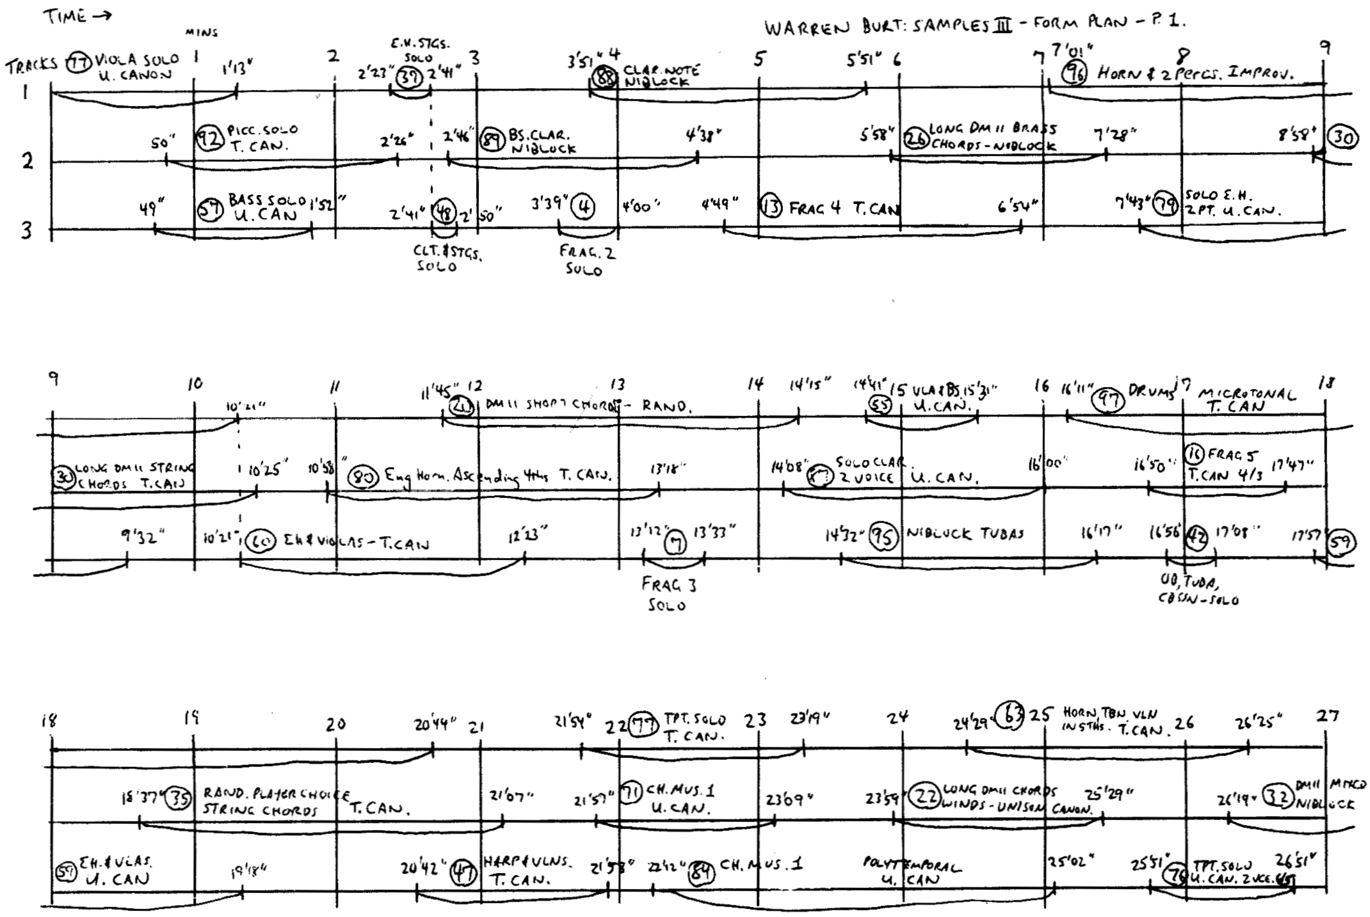

Ex 2. Structural outline of the first 27 minutes of SAMPLES III.

beating, etc. is really an artifact of the electronic process, and how much could be obtained under other acoustic situations, such as two or more instruments in a room playing notes tuned extremely close together.

3) Random textures obtained by having Aardvarks IV "play" the Akai, on which were loaded a number of short samples, each of which could be played only over a pitch range that did not disturb its "timbral identity."

4) Improvisations made with either single notes or longer samples to assemble complex textures, such as the trombone glissando sections or some of the percussion and gong textures. 5) Playing of single samples unmodified and unmixed into the final mixing process. This was mainly done with each of the full orchestral and some of the chamber textures.

All of this processing gave me 100 recordings, ranging in length from 8 seconds to just over 4 and 1/2 minutes. Harmonically, they covered four areas- diatonic harmonies of various types, chromatic and random harmonies using the full range of chromatic pitches, microtonal textures, and noisebands resulting from either extremely thick textures, or from non-pitched percussion. Needless to say, on some occasions, the boundaries between these categories were considerably blurred. This work was accomplished in September 87.

The 100 sections were then combined and mixed into the final 84 minute form of the piece. Again, the idea was to have sometimes one, sometimes two, and sometimes a maximum of three textures playing simultaneously. For convenience of tape lengths, it was decided to divide the piece into two sections of roughly 43 and 41 minutes each, though in radio broadcast or public performance, the two sections would follow each other with only a couple seconds silence between them. In keeping with the rest of the structure of the piece, the overall form was also determined with the aid of random processes. In this case, *A Million Random Digits* (a book chock-full of chuckles, I can assure you), was consulted to first determine the number of recordings that would be present in each track, and then to select which recordings those would be. The total durations of those recordings were added up, and the result subtracted from the total length of the section, and that duration (the amount of silence on that track) was divided by the number of silences between the recordings. This gave a duration by which the recordings were separated on that track in the first sketch of the overall form. This resulted in a three part texture with considerable silences, as I was aware it would. I now had the choice of leaving the silences as they were, or of modifying them in some way. As a sort of parody-homage to radio technicians' fears of "dead air" (silence), I decided to eliminate silence from the piece entirely. (The loving spirit of John Cage is hovering around enough of this piece already without also buying into that one!) So any silence that resulted from this structure was eliminated by sliding the locations of my fragments backwards and forwards until the entire texture had a "neat fit," either through overlappings, or several times, through simple successions of textures.

Example 2 shows the structure of thie piece for

its first 27 minutes. *U. Can.* means the texture is an untransposed canon of some sort. *T. Can.* refers to a transposition canon. *Niblock* means a microtonal texture of some sort. *Solo* refers to an unmodified fragment mixed in. *Improv.* refers to a section made by improvising with samples at the keyboard. *Rand.* means a section made with Aardvarks IV controlling the Akai. Numbers in circles refer to the numbering system used to select the 100 fragments. It will be seen from this how the random process resulted in a form which juxtaposes various kinds of textures. For the mix, levels on the three tracks were kept about as equal as possible. Occasionally, if a texture was completely obscured by another, some adjustment would be made. Overall, the aim was to get as transparent a mix as possible, so that various textures could be "heard through" each other when they were juxtaposed. Domination of one sound by another was not usually desired. Rather, a sense of being able to appreciate the coexistence of very different sonic worlds was what was usually aimed at. All three tracks were mixed to stereo, with track one being slightly to the left, track two centered, and track three being mixed slightly to the right of the spread.

The piece was broadcast on 17 November, 1987 as part of ABC-FM's late lamented *Audio Spectrum* program. The media greeted the appearance of a new work for radio with its usual complete silence. Slowly, however, various friends reported hearing the piece, and others listened to it on tape, so that response to the piece began to slowly filter through. Generally, this response has been favourable. I would like to hope that this long, thick and difficult piece would be one which people would want to come back to time and again, to explore the many levels contained within it, to use as a means of exploring their response to sound, to media, to the orchestra, to orchestral recording, to extreme duration, to various kinds of "difficult" textures and juxtapositions, to explore both as a "thing in itself" and as an element in a rich and complex web of reference and interplay.

However, I would like to reaffirm, that for me, this piece is a radio piece and a tape piece with its own distinct identity as such. It is not a piece where the orchestra has been "appropriated," nor is it simply a piece for electronic sounds. Rather it is one which uses sounds I like in a way which fascinates me. That is, it explores a timbral world which interests me with a physical medium of sound production (notice I did not say reproduction) I find particularly intimate and appealing. This point might be expanded a little. In much contemporary visual arts criticism, one finds easel painting referred to as the "signifier of cultural authority" which artists are now "reappropriating" after years of other kinds of activity. It might be viewed that my use of the orchestra in this piece constitutes just one such kind of "reappropriation" of a sound-world of "cultural authority." But I don't feel that this is at all the case in this work. First of all, "cultural authorities" exist not in societies, but in the minds of them what gives authority to them. If after all the work we have done over the last 20-30 years establishing a variety of media each with their own unique qualities, we still feel that any one form has more "authority" than another, then I think we've done a lousy job of establishing the uniqueness and dignity of these many new forms. Further, I feel that it might be time to call into question the dominant metaphor behind much contemporary critical language. Look at the words that are used again and again by contemporary critics: "appropriation," "incisions," "violations of codes," "complicities," etc. To anyone with half an ear for language, this will clearly seem like a language, and a particularly romantic language at that, of war, subversion, and violence. The military metaphor is the dominant one used in critical discourse today, with critics romantically viewing the artist as either a subversive or a compliant force within a repressive society. Well, society IS as repressive as our critics say. There's no arguing with that. However, if we want to build a non-violent society, one based on peace, co-operation and trust, perhaps its time we started working on ourselves, and developed critical metaphors that were not related to war, spying, and violence. To use a non violent language that is fully reflective of our desires for non-violence- to use a language which does not participate in the heirarchical structures that the art it is describing attempts to provide an alternative to. For example, my broadcasting of this 84 minute piece on ABC-FM could be described as "appropriating media space to deconstruct the codes of duration!" This description, though acknowledging that most media structures impose on us a very short attention spans, ignores the fact that classical music radio stations are used to broadcasting works of long duration. Mahler, Bruckner, Glass and Ashley's longer works are semi-regular features on ABC-FM. Furthermore, my presence on that station was as part of an officially sanctioned project. If there is anything revolutionary about the work and its placement (and I feel there is), it will be realised as a result of people dealing with it seriously, seeing what is implied by both its unique nature, and by the structures it embodies, rather than by simply describing its being broadcast in facile and overly romantic terms.

And this kind of implication of theft and violence exists right throughout our language. "Synthesiser" has come to imply that somehow, the sounds made by electronic waveforms are "fake," and "unreal." "Sampler" has come to imply theft or "shoddy goods," as in the peddler with his "samples." As anyone who has heard a square wave at 130 decibels can assure you, there is nothing "fake" or "unreal" about simple waveforms. And there is no inherent theft in making a recording, digital or otherwise, unless one chooses to use it in that way. That these instruments have the potential of changing their identities in a way other instruments don't is true, but to think of them as somehow "bogus" or "thieving," and to think of works on tape as being somehow "unreal," strikes me as illogical, a little like referring to a computer as the "thief of mathematics."

I prefer another implication of "sampler," one derived from the non-violent, traditionally female art of quilting. Like those "samplers"- made from many small bits of cloth, some found, and some specially made- I would hope my "Samples" would serve a similar function- as useful, honest, and homely objects for people to engage, to work with, and to enjoy.

Sounding and sampling

Contracted notes on the conceptualisation of music technology

Philip Brophy

The history of music technology in this century is generally presented as a series of interfaces between man and machine, coloured by man's endeavours and the wonders of machines. Rooted in the programmes, forecasts and desires of the industrial revolution (and therefore slightly suspect in a post-industrial epoch) the predominance of the man/machine interface indicates ways in which music technology is still conceptualised. It is not surprising, then, that the more developed and sophisticated music technology becomes, the more generalised and schematised that interface, resulting in pseudo-neutralised terms like 'user' and 'hardware'. But while there has been an undeniable increase in music technology applications over the past decade (especially with the merger of professional and domestic usage) which to an extent accounts for this symptomatic streamlining of the man/machine interface, it would be unwise to view historical antecedents and formations of music technology as somehow 'inevitably' leading us up to the digital revolution. There are other possible ways of viewing the history of music technology. This short article suggests some, by focusing on how the **concept** of a man/machine interface has informed many major developments (technical and theoretical) in music technology.

Accepting that the very concept of such an interface is historically and culturally determined (transplanting renaissance man into the machine age to foster a romance with technology) and therefore subject to change through time, there is no reason why one can't now invent different projections and readings for 'man' and 'machine' in regards to the subject of music technology: if one can view 'man' as the active producer of sound, and 'machine' as anything that can be sounded, it then becomes difficult to imagine how an interface between man and machine could **not** eventuate in the event of a sounding. Consider these two examples: (i) man deliberately controlling his breath rate to produce a tempo and rhythm to be experienced **as such** (man becomes machine): and (ii) man manipulating dead matter (hollow log, stretched hide, etc.) to make objects for the production of sound (machine needs man). One wonders if man's musical existence could at all be separated from his technological (physical/mechanical) relation to the (internal/external) world.

These two examples (many more could be projected) are presented to suggest that the man/machine interface is not a construct but a **given** — a paradigm that doesn't need to be demonstrated because there is little chance of escaping it. This means that a history of music technology which establishes and promotes a framework of man/machine interfaces (pinpointed chronically and mapped out chronologically) is an inept model for articulating the man/machine 'abstract' — that is, the ways in which man generates and articulates a musical existence by producing and working with technologies. The implication here — a feasible one, too — is that every instance of man's musical existence is in fact a juncture of the musical and the technological; an event predicated on the situation of actively 'sounding' something.

The man/machine interface would be better conceptualised as a **matrix**: a self-interlacing and self-superimposing configuration of interactions, dialogues, dialectics, usages, manipulations, abuses and exploitations. Here the man/machine interface is broken down and detailed as a fractal network constructed by and between (in no order) man, machine, human, device, composer, technology, moron, instrument, artist, material, musician, sound, creator, tool, etc. etc. etc. The angle of connection (as an **incident**) is more important then the interface itself. Furthermore, this network is not constructed: it **eventuates**, as lines drawn out and across every juncture of the musical and the technological. These junctures could be termed 'ordinances': events upon those projected and directed lines which shape this network's expansive and pervasive matricular form.

If there is a 'logic' to this spread of matrices it is to found in the ceaseless **collapse** between man and machine, wherein each musical/technological event creates its own ordinance in the matrix, virtually independent of man's endeavours and the wonder of machines. A 'network of matrices' is thus conceptually more attuned to the continual flow of occurrences ('soundings'), for each ordinance is not simply a musical/technological juncture, but more precisely an instance of this col-

lapse. To boot, each event of 'actively sounding something' fails to distinguish between the two most conventional visions of the man/machine interface — invention and application. To invent is to apply, and to apply is to invent, especially in regards to man's perception of himself as a machine and the machine as a displaced self. This accounts for the formation of projected and directed lines in the matricular network — maintained, customised and reassembled as they are by each and every application/invention of music technology.

A key figure in this blurring between application and invention is John Cage's prepared piano. It is simultaneously a renovation of the object and a reinvention of the instrument. But still, one is left to ask: prepared for **what**? Countering the sonic tactility of this instrument, the sounding of the prepared piano smacks of the wonder of transformation, where each tink and clunk declares its refined 'nonpianoness'. Historically important as it is, this invention/application is typically ignorant of the predestined collapse of such human interfaces and modernist tamperings. If one considers the strategy of deadening the strings, constricting their pitch and recontrolling their hammers, one is actually brought back to the original design principles of the piano — indeed, of all musical instruments made of dead matter and sounded through acts of violence (pressure, force, friction) by man. Cage can yen his way until the globe becomes flat, but he is no less necrophiliac than the deadest of European culture and its morbid 'mastering of nature'. He simply prepared — or mummified — the piano differently from its original process, throwing nuts and bolts into the works like an anarchist tampering with the machine to signify the **noise of pitch**. Such are its most interesting points.

The prepared piano perversely prepares us for nothing — especially if one accepts that musical/technological distinctions and frameworks based on man/machine interfaces and constructs were and are illusions in the first place. The prepared piano is essentially a declaration of transformation: a modernist gesture of sounding which calls attention to the act more than its sound. In a sense, an 'empty' gesture, but an emptiness that is in the nature of the **flow** within the network of matrices of music technology developments which **empty man and machine of each other**. The integral impulse behind such developments is not only the desire for invention (of the new) but a mania for **supersession** (by the new). Obsolescence is not planned as much as it is readily accepted. Each 'invention' (presented as a development) is thus ultimately an empty gesture, but nonetheless capable of generating lines of direction and projection for the supersessional flow. This is all relative to the matricular network wherein everything is to be continually adapted and replaced — an activity perfectly suited to the modernist brief of building upon the past.

While the prepared piano adapted an object (the piano) and replaced its sound (with 'nonpianoness') its status as gesture constitutes an attempted blockage in the flow of invention and supersession — to halt things in order to draw attention to the act. Of course no such blockage occurred: the prepared piano did not **become** an instrument of its own making. One may well ask today (rephrasing Barthes' *Musica Practica*) "Who plays the prepared piano today?" The 'reason' for its existence is primarily artistic/philosophical and not necessarily utilitarian. Yet as an artistic endeavour the prepared piano paradoxically rides the conveyor belt of artistic **needs**: Cage 'needing' to break down musical barriers in order to explore sonic possibilities; Varese 'needing' the tape recorder 20 years before it was invented; Schoenberg 'needing' a means for the emancipation of dissonance; Partch 'needing' new instruments for his musical sound; etc. Popular yet suspicious claims. Artists simply feel the need to create, to invent, to produce; some might even call it a neurotic condition. Their drives are manifested in their objects and compositions, marking them as true inheritors of the 'empty' impulse to invent and supersede, to continually maintain, customise and reassemble. Synchronous with the creative impulse in 20th century art, true obsolescence in music technology is to be found in the re-exploration and reinvention of areas which already are adequately serviced by an existing stage or phase of development. Still, empty eventfulness of the **non-invention** and **re-design** of instruments and machines is central to the man/machine matrix, where man feels a need to invent, create or produce something when there is no extant use for it. The 'need' is created in the event of its creation.

I am not denigrating the artistic value of the creative impulse. **I am** questioning the use value of technological invention by attempting to demonstrate its links with the core vacuousness of creative activity. There is nothing perverse about this if one believes that art and technology are the two most amoral practices of the 20th century (as vociferously proposed by the dadaists at the start of the century). To discuss these practices under moral terms thus to me seems imperceptive or dismissive of the subsequent delusory spread of these practices' purported ethics and ideals which — under their banner of creation and invention — are solipsistic at heart. If one reconceptualises the man/machine interface — in a more amoral and less humanist tone, and with more **flow** and less **structure** — one could better realise alternative performances and functions of music technology : the collapse between man and machine; the lack of distinction between invention and application; the inevitability of warranted and unwarranted supercession; the neurotic impulse to adapt and replace; and the act and event of non-invention.

From here we move on to what is perhaps the key contemporary issue in music technology, especially in a socio-cultural context: sampling. While various binaries and dichotomies have set sampling up in terms of simulation versus representation, such a critical method ignores the greater historical lineage of music technology (which I have attempted to introduce as matricular in form). If 'sounding' is the situation of actively producing sound, I fail to see how sampling — as a mode of production — can be severed from all other means of man manipulating materials to produce sound. I can (barely) comprehend a certain folky melancholy, or even a dogmatic ideological stance, both of which express concern over the 'direction' in which sampling is headed, but such

concerns put preformed fears of the state and condition of sampling before considered observations on its performance and function. In reference to the afore-mentioned conceptualisation of the man/machine interface, I wish to suggest that this 'issue' of sampling versus sounding is another collapse; another ordinance, another collision point to be found in the man/machine matrix. But for it to be accepted as such we need to trace some of those bad binaries (whose progressivist/positivist lineage is beyond the scope of this article) that posit digital electronics as some sort of new and fearful dimension in music technology.

We start then with analogue electronics (often termed 'reductive' in nature) where **filtering** is the key operation. To filter is not simply to reduce: it involves selecting, shading, shaping a sound. In effect it is a form of 'culturing' a sound; of transforming frequency to pitch and back again through voltage control; of demonstrating the degree of manipulation that determines the act of sounding. The Moog synthesiser — as the most famous instance in the domestication or 'musicalisation' of electronic synthesis — based its design/invention and adaptation/replacement on transforming the keyboard. Here filtering was put to the prime service of generating a change in tone corresponding to a rise in frequency (such as when shorter lengths or greater tension in string vibration and wind passage flow cause a sharper tonal definition). Seemingly 'naturalistic' in its drive to replicate a pre-existing set of acoustic musical technologies, this mode of filtering also draws attention to itself, to its feat in **controlling** sound in the act of sounding. This is the prepared piano revisited — this time **removing** the nuts and bolts, the tinks and clunks in an attempt to naturalise the electronic, to effect not the noise of pitch but the **sounding of pitch**.

Most importantly, this desire or tendency to filter is determined by a particular conceptualisation of the keyboard as an operating base or general headquarters for an envisaged man/machine interface — indeed, the keyboard is a construct of such an abstract. The concept here is of a keyboard as a flow chart, facilitating a series or sequence of pitches; a (pre)arranged and controlled flow (left to right; low to high) of the key elements in musical composition. Essentially, the relative change in tone through filtering is to reaffirm the original design of the keyboard as a device that maps out pitch, because the keyboard now maps out tonal change as well. The incorporation of a keyboard into sampling components evidences a very different notion of the keyboard. In sampling keyboards there is no flow or sequence, but rather a **continuum of breakages** where each key is an entity — an isolated trigger — which only happens to be relative to notions of pitch. Each key or note on the sampler keyboard is not only total as an event (the capture, snare or snatch of a sounding) but also in material, substance and form. Quite simply, this means that as you move up and down the keyboard, the sample (or 'sonic entity') is accordingly altered in pitch, duration and timbre — **replicating** the contrasts effected by magnetic tape speed changes. Each different key struck thus makes one aware that one is experiencing a **transformation** of an 'original' sound. Yet another visitation by the prepared piano and its mode of transforming a pre-existing sound identity. Ironically, it is the state of too much filtering or tonal change which renders the sample keyboard artificial.

The point here is that while Cage's prepared piano signaled the end of the keyboard academy and the start of the experimental apparatus, the re-employment of the keyboard in analogue and digital sound generation has largely determined both the dead-ends and new horizons to be reached in this realm of music technology, especially in regards to the merger of professional and domestic usage. To fully appreciate this one only has to consider (in logical order): (a) the reaction of musicians against using a computer keyboard in digital synthesis and musical composition; (b) the number of parameters in analogue and digital synthesis geared around keyboard manipulation, interaction and performance (sensitivity, dynamics, tonal contrast, etc.); and (c) the rampant/prolific reconstitution of sampled sounds into musical soundings, where a 'pattern' played on the keyboard transforms the isolated sample into a melody. Cage was one of the first to seriously (albeit rhetorically) question "Who plays the piano today?" Well, **today** a legitimate answer confronts us: **anyone and everyone.**

In sampling applications the collapse of the man/machine interface is further marked by an open disregard for who is capable of what, giving us a socio-cultural breakdown between composer and listener as opposed to the polemical/rhetorical claims of experimental music. This is to be found in the dissolution of the roles of producer and consumer, for with sampling, to produce is to consume and to consume is to produce — all at the push of a button, the trigger of a key. While some bemoan this to be some sort of frightful last straw (based on musical ethics) the point is that the production/consumption dissolution is far from being a recent phenomenon.

One could easily view the diatonic scale — architecturally spread across the keyboard as set of dos and don'ts (sharps, flats, naturals) — as the presentation of a predetermined frequency range of **spectral excerpts**, designed for consumption more than production. To his credit, Cage's prepared piano was intent on experiencing a fuller spectrum than the keyboard's excerpts allowed. In this sense, Schoenberg's emancipation of dissonance **via** the keyboard is but a conceptual victory, leaving others like Partch, Ligeti, Bartok and Penderecki to reinstate frequency and flow into the architectonics of musical composition (by sliding, blurring and dissolving pitch rather than constructing it). To put it another way: Cage tried to position himself between the strings of the piano frame; Schoenberg tried to move unconsciously between the lines and spaces of the music stave; and Penderecki et al tried to disappear into the cracks between the keys on the keyboard. In this light, 20th century music is as much about the keyboard as it is about tonality, harmony, sound and acoustics. Not surprisingly, the keyboard is prominent in music technology: whereas both acoustic and electronic keyboards control frequency as pitch, the sample keyboard controls sound as pitch. And just as sampling is in a sense the co-

option and corruption of musique concrete's material manipulation, it is likely that all keyboard music — no matter how avant garde or experimental — in some way requires the constriction and control of sound (through filtering, selection and excerpting), and that the keyboard (as abstract and construct) is largely responsible for such a condition.

Finally, the sample keyboard signals the return of sounding through dead matter. It is a 'return' because analogue electronic synthesis marked the first crucial shift from sounding the sonic potential in dead matter to the sonic application of live energy — the **sound** of electricity. Based on voltage control, it transforms live energy into material for musical/sonic manipulation. Analogue synthesis is thus mostly 'eventful' in that one is **continually** controlling energy in order to shape a sound, as opposed to acoustic sound production which is determined by the employment of an instrument which is **already** a crafted and completed object designed for the production of sound. The application in the former is live; in the latter it is given; electronic sound is generated while acoustic sound is produced. Digital synthesis — with its modulation of frequency (FM) — then marks the first shift away from the totality and unification of mediated/unmediated energy, dealing with a multiplicity of and interaction between energies (the principle of 'additive' synthesis). Sampling devices and systems effectively confuse these issues of energy control, because energy is **subsumed** into the servicing of record/sample functions. While the sounding of a sample is dependent on an internal (microchip) control of energy, the **sound of sampling** signifies the absence — or phantom presence — of energy; of either an acoustic occurrence, a material transformation or a musical event. Samples **sound recorded**; in their simulation of the real they only confirm their artificiality and prove their illusion.

The 'dead matter' this time is thus not the physical state of the instrument, but the nature of the material, substance and form of the sample. Here, issues of simulation and representation are overidden by notions of **animation** and **reanimation**. The point though is the **return** to and of dead matter, for just as the act of acoustically sounding something 'animates' dead matter (causing particle resonance and effecting wave displacement to produce vibration as sound) the sample 'reanimates' such an occurrence by absenting matter in the very act of sounding (disintergrating the difference between sounding and sampling). While sampling causes confusion due to its hyperreal effect which is amazingly not dependent on real material and matter, sounding and sampling are remarkably similar to one another, for in the end, both necessitate an act of sounding. This is the nature of their collapse into one another. Both means of sound production testify not to the death of silence (a physiological impossibility) but to **silence's inability to sound us** - its inability to make us conscious of sonic potential **without** being put in the situation of actively sounding something, wherein our musical existence is determined by our technological (physical/mechanical) relation to the (internal/external) world.

Towards a virtual piano action

A report on an Artists and New Technology project[1]

Alistair Riddell

Piano Music in Evolution

Since the turn of the century music for the piano has exploded with stylistic diversity and performance innovation. This has largely taken place where the inherent qualities of the instrument and their potential have been recognised and promoted. Not only in the west but in countries such as Japan, China and Korea, enthusiasm is manifest through performance, manufacturing and the fostering of study and social awareness through music schools and methodologies (Yamaha, Suzuki, etc). The instrument's historical repertoire (essentially from the Viennese classics to early Twentieth century) continues to feature prominent-

ly in most assessments of the instrument. However, it is clear that through many unique social and technological changes, its musical potential is being exploited by a far wider and larger cross section of world societies. This is evident, for example, in the scope of discussion in popular keyboard magazines. It may also be assumed that the process is continuing under the influences of today's technology.

In certain musical idioms, particularly those in the aural tradition such as Jazz, Blues, Rock, etc, piano technique has changed through the dissemination of recordings as well as through direct contact with the current influential practitioners. The dissemination of recorded music has supported many enthusiasts who, by emulating the definitive performers, have themselves developed further distinctive techniques. In a genre of music that appears to be less constrained by formalism and pedagogy, imitation becomes an underlying reactionary force that stems the tendency for particular styles to disintegrate as quickly as they emerge.

The situation in the Art music world — as distinct from other musical milieus: commercial, popular, etc — is somewhat different because most performers are bound to interpret some form of score. In this notated representation of the music, technical as well as stylistic innovation is dependent on the composer's practical and theoretical knowledge of the instrument. Thus, the success of music composed in this manner is largely attributable to the perspicacity of the composer in the treatment of both instrument and performer. This is evident in such works as the prepared piano works of John Cage; the *Klavierstucke* works of Karlheinz Stockhausen; Charles Ive's *Concorde Sonata*; Elliott Carter's *Piano Sonata* of 1945/46; Pierre Boulez's *Structures*; the Player Piano *Studies* of Conlon Nancarrow and many others from the twentieth century repertoire.

Recorded music is perhaps less influential as a conveyance for imitation in an art music context because direct imitation is not explicitly part of the evolution of that musical genre. However, recordings are evidence of the success of a musical idea and convey proof that certain compositional ideas work (or don't). Furthermore, the quality of contemporary music performance and interpretation has improved through its general availability in a recorded form. Those interested in contemporary music are able to further familiarise themselves with it between performances. Recordings also tend to establish personal criteria through which a comparison can be made between the memory of the recordings and live performances or other recordings. The ease with which it is possible to compare one interpretation of a work with another or one style with another, must contribute to a broadening of musical perspective and at least, perhaps introduces avenues for change.

Technical Evolution

Underlying the piano's social ethos throughout this century is the fact that the instrument had reached evolutionary stasis by the end of the nineteenth century. The following is a summary of the significant developments of last century (Grover 1973. p.210.):

1821. Sebastien Erard produced the 'repetition action' with double escapement.
1826. Jean-Henri Pape patented felt hammers.
1843. Jonas Chickering patented the first one-piece cast-iron frame for a grand piano.
1859. Steinway produced the first over-strung grand piano.
1874. Steinway perfected the 'sostenuto' pedal.

Certainly, since last century the piano has undergone considerable refinement in component materials and production techniques yet it has essentially remained an instrument of 19th century technical achievement. It is perhaps, possible to attribute a considerable part of the diversity of piano music this century to developmental stability. Unlike the current state of electronic musical instruments, musical innovation is entirely up to the performer; the manufacturer has no part in encouraging new music other than through promoting artists who use their instruments.

This illusion of stability should not be interpreted in the first instance as an attainment of technical perfection. Given human nature and art, this is, of course, intrinsically impossible. The instrument, however, remains relatively constant in appearance and sound because of its social prominence and musical tradition. In order to remain a worthy interpreter of the historical repertoire as well as the piano music of tomorrow, the impact of even minor technical modifications and improvements are generally vetted against the standard repertoire, performance criteria and contemporary musical expectations. Transitions from wood or ivory to plastic, or from iron to alloys, do not significantly affect the perpetuity of the piano repertoire and are eventually accepted without regret. Yet the sum of the many minor changes must ultimately be manifest as an evolutionary step, even if obscured by the passing of many years.

Although it is difficult to clearly define the relationship between instrument evolution and musical style, music itself should not be assessed on the basis of technical progress. As Willi Apel pointed out:

> "Nothing is more dangerous and misleading in the study of the arts than to regard achievements of the past from the standpoint of technical progress. A superficial observer sees only what has been gained in the fight and not what has been lost. The true historical mind, however, sees that in the history of humanity there is no possibility of perfection, and that there is only a faint hope of approaching it." (Apel 1953. p.86)

To some extent the continuing popularity of the piano repertoire of the 18th and 19th centuries is attributable to the fact that its performance has been widely accepted on contemporary instruments. Those with an understanding of the historical instruments would argue that certain original qualities of the music are consequently lost or misrepresented. This may very well be true but for the most part the musical intention and aesthetics are present where the interpretation is sincere and reverent.

This concern for historical accuracy has increased with the growth in popularity and scholarship of pre-baroque music this century. It has fuelled the underlying controversy and debate over, not only

the use of authentic replications of historical instruments but musical practice and interpretation researched from primary sources. Stepping back from the issues surrounding early musicology, the net effect of this debate is perhaps to increase the musical public's awareness of instrumental diversity and evolution. If there is greater awareness of a musical continuity then contemporary instrumental practices should appear less disconnected from tradition.

The technical evolution of musical instruments is, of course, inevitable. What is perhaps different today is the nature of the association that is formed between the various instruments and the current musical styles. As the instruments themselves become 'virtual', that is able to adopt the characteristic sounds and behaviour of many other instruments, the relationship between instrument and musical idea that is not only difficult to define but transient. Combine this with the rate with which 'new' instruments are appearing on the market and the result is a instrument/music relationship that has never existed before.

The Action Project in abstract

In 1987 I began to consider the design and development of a high performance electro-mechanical action for a grand piano. The opportunity arose through the Artists and New Technology Program from the Australia Council. The concept had been maturing from the early 1980's when I first began working with pianos under microcomputer control. During those interim years, composition and research with less ambitious instruments and equipment (Riddell 1982, 1988) helped clarify my future intentions, aesthetics and the inevitable pragmatic considerations associated with the project.

When construction seemed imminent, expectations of the function of the action were largely influenced by the prospect of an unusual Performer/Machine Interactive system. One that could exploit a distinct relationship between the piano tradition and digital technology. At that time it was felt that mainstream music technology was not pursuing the same type of instrumental relationship in such a simple and direct manner. On a fundamental level, the difference is manifestly the acoustic production sound, while at a higher level, it is performance on a familiar instrument with an entirely different philosophy towards its control.

An integral part of the action's functionality allows it to be used in two performance contexts. In the first, it can operate through a live performer in real-time. In the second, by the computer alone, that is without any real-time human intervention. This range of functionality covered my interests in real-time performance and algorithmic composition.

From the perspective of a performer/machine interactive system, the action can be regarded as 'virtual' in that the performer's actions can be mapped to almost any pitch/rhythm combination on the instrument. Transposition, inversion and intervalic parallelism with all or partial input are some of the simple processing functions possible for keyboard actions. The more fascinating possibilities lie beyond these techniques where performer input is mapped to parameters other than pitch or rhythm. In fact they might not be parameters but functions.

The action has also a degree of operational autonomy that the traditional action could not possibly accommodate. The relationship between a hammer and its associated damper is flexible. They can operate synchronously or under separate control which permits two timbral possibilities: playing with the damper off or playing with the damper on the string.

Although the prospects for the research and performance of the action appear limited, it is an attempt to alleviate the difficulties of working with computer technology and acoustic instruments. To begin with the grand piano is a complex and yet ideal medium that is reasonably available at most performance venues. With the action in a modular state, the logistics of performing are considerably reduced. However, installing and adjusting the action will still take some time but does not require the piano to be altered or modified in any way other than the removal of the original action. This is a straightforward operation which takes only a few minutes.

Initial Research and Observations of the Traditional Grand Piano Action and the Proposed Action

The preliminary research towards the project began in mid 1987 in the Department of Physics at La Trobe University[2] This initial research was not directed towards the immediate construction of the new action but an examination of the physical behaviour of the traditional grand piano action. The research was conducted jointly with Dr Michael Podlesak[3] and centred around the measurement and investigation of the velocity, acceleration and power in a hammer when propelled into motion. Experiments were carried out using an ONO SOKKI FFT (Fast Fourier Transform) analyser. This instrument recorded, calibrated and plotted values that were transmitted to it from a tiny accelerometer mounted on the hammer shank.

The information gathered from the experiments was used to convert the hammer's mechanical energy into electrical energy. The power in the hammer motion was quantifiable in terms of Watts and these values could be used comparatively with the electronic solenoids also under examination.

The results finally revealed that, within certain limitations, the physical size and power rating of the proposed hammer solenoids would be acceptable. This was particularly important because it meant that the existing Pianocorder solenoids and control components could be used, thus reducing the cost and complication of the overall system. The major restrictions, however, arose in relation to the adequacy of the power supply from a standard domestic source (2400 - 3000 watts) and whether the solenoids would be capable of an effective dynamic range in considerably less than optimal conditions.

The experiments on the instrument focused on two dynamic ranges: moderately loud and extremely loud and the differences between human and electro/mechanical performance for this experimental range was considered from a number of perspectives.

The more demanding end of the performance spectrum was of more interest since the lower end

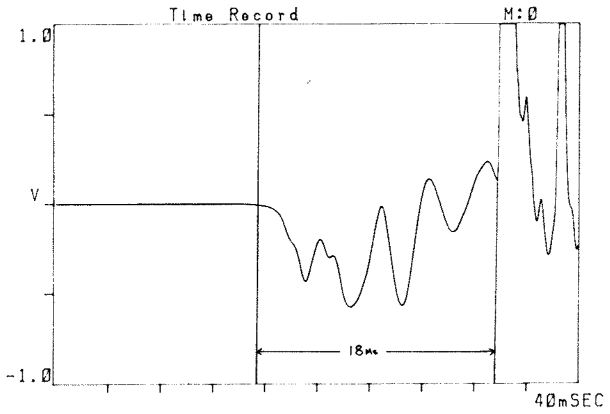

Figure 1. An FFT plot of the hammer striking the lowest A on the grand piano.

is demonstrably less problematic for either system, attention could therefore be directed towards the extreme case to at least establish a possible upper limit.

The experiments at this dynamic level revealed that the power in terms of electrical energy was approximately equivalent to 25 watts. However, the occasions for such extreme performance (if any) are only likely to be found in later twentieth century music and even then, in order to achieve such a dynamic, the performer would generally have to forgo speed, elegance and repetition for brute strength. It might be expected that even the most extreme contemporary works would not require the performer to sustain or distribute this sort of power to large groups of notes over a considerable time. Neither is it possible to apply such power to complex passages with any degree of subtlety or accuracy (although the Jazz musician Cecil Taylor might come close to being an exception).

It has been speculated that the experimental solenoids could take an equivalent power (25 watts) and possibly deliver the same effect. As desirable as it may be to emulate this level of performance, it is doubtful that the existing solenoids and support electronics could sustain repeated use under such conditions and also whether many of them could be supplied from a domestic source. This result was eventually considered too extreme to pursue as necessary for the electro-mechanical system yet it helped to define a practical upper limit.

The moderately loud attack resulted in power ratings between approximately 7 and 10 watts. This considerably reduced range was closer to estimated realistic performance figures for the solenoids. The solenoid experiments were carried out with relatively conservative power supplies and it was speculated that their performance may be improved through more efficient and large power supply strategies.

The issue of power supply and usage was further complicated by the differences between the mechanical operation of two hammer systems — conventional and electro-mechanical. The hammers for the electro-mechanical arrangement (discussed in the next section) are heavier than the heaviest piano hammer by about 7 grams. The heaviest bass hammers on the conventional action are approximately 13 grams while the iron solenoid cores alone (without the actual hammer tip) weigh 20 grams each. However, a larger hammer mass with less velocity could result in a similar dynamic to a conventional lighter hammer travelling at a

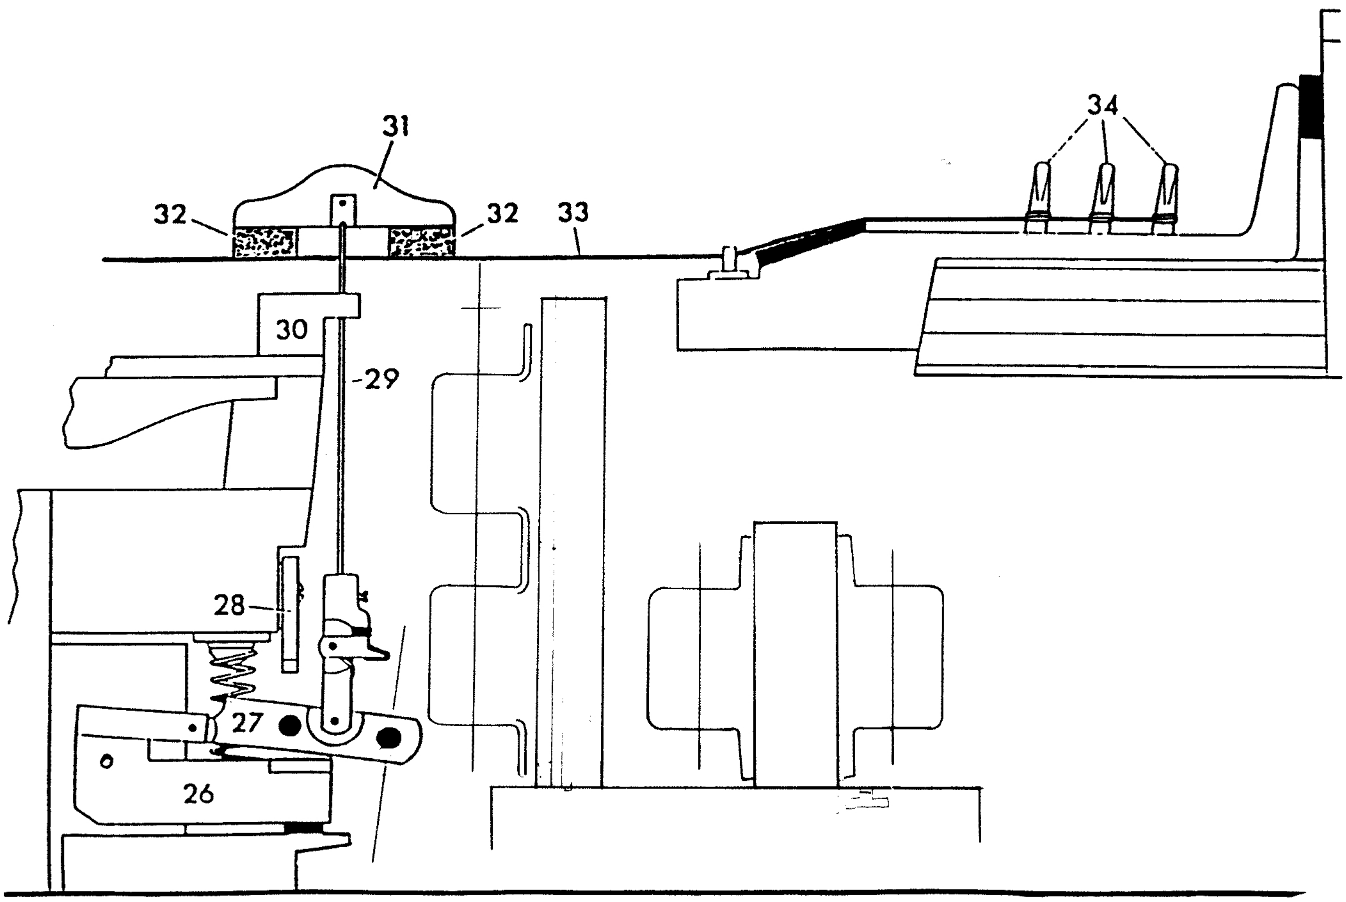

Figure 2. A scale drawing showing space restrictions inside the piano.

greater velocity. This was not experimented with directly but was recognised as an important factor in the performance of the new action.

In addition to the weight differences the operation of the two mechanisms necessitates an alternate assessment of the mechanics of the hammer movements.

The FFT plot (Figure 1) shows that the time taken by the hammer from initial impulse to string impact was approximately 18 milliseconds (ms) for the moderately loud case. From that time onwards the hammer is returning to a rest position. The further oscillations recorded on the plot are the movement of the hammer on its wooden shank. For the initial thrust of the hammer the plot does not show the point where it is free from the impulse of the keystroke force and travelling without further propulsion. Although this period of free flight may well be negligibly small, by contrast the solenoid remains under power and acceleration until it impacts on the string and potentially interferes with its vibration.

It is expected that the solenoid can be powered up for an approximately equivalent period of time as the normal hammer — about 18ms. But that also depends on the intended dynamic. The dynamic range for the electro-mechanical action is achieved by turning the solenoid off at various stages after activation. The loudest attack will result from the solenoid remaining powered up fractionally beyond the time of impact. This may be 20ms while the softest attack might only be a 6-7ms period. An influential factor in the timbral quality of a struck string is the reflected waves from the shortest end of the instrument which return in approximately 9ms. The longer the solenoid remains in contact with the string, the more the hammer interferes with the evolving spectra.

Construction of a New Action

The question of performance was momentarily put aside when time came to consider construction of the action. After all, abstract performance questions were of little consequence if it turned out to be impossible or impractical to construct a mechanism — using the resources at hand — that could fit into the complex cavity that normally houses the conventional action.

Through the theoretical principles of operation and those existing components that were inextricably part of the 'grand design', some idea of its form was known in advance. However, from February of 1988 to late June, many ideas and approaches were examined and discarded as the search for a workable design was given priority.

The action began to take shape in the hands of Marshall Maclean[3] in the Physics workshop at La Trobe University from June of 1988 onwards. Many of the technical and design problems that sprang up during construction were resolved by him (see the working sketches). Without recourse to a piano at every critical point, he produced an action that constantly met not only with theoretical expectations but proved successful during intermediary tests.

The action benefited from many important decisions made along the way. Some of those were:

- No specialised solenoids were to be constructed either for the hammers or dampers. Those used in the Pianocorder mechanism were considered to be suitable provided they could perform to an acceptable level and be located in the already cramped space. Development went ahead on those premises.
- The action could be made to fit a number of instruments by being modular and adjustable. The major obstacle, however, that this attempted to overcome in the first instance was the intense asymmetry and irregularity of the piano's construction. It could not be assumed that internal structure, for example, would remain constant or that it would change in a linear fashion. Consequently, consideration was given to the possibility that part of the action mount might cause interference with the action frame, i.e. frame struts or sostenuto bar.
- A preliminary test installation revealed that there was more space available than initially thought. This meant that the action could now reside well within the limits and permitted greater variation between instruments.

An examination of the photographs of the first (bass end) of the three modular sections identifies the support structure for the hammer and the damper solenoids. These are arranged in two parallel rails mounted on adjustable feet. The hammer/damper configuration met the functional

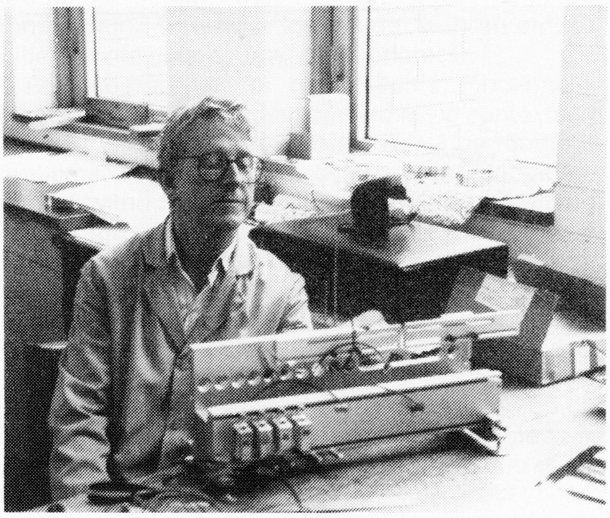

Marshall Maclean with the first of the three sections of the action.

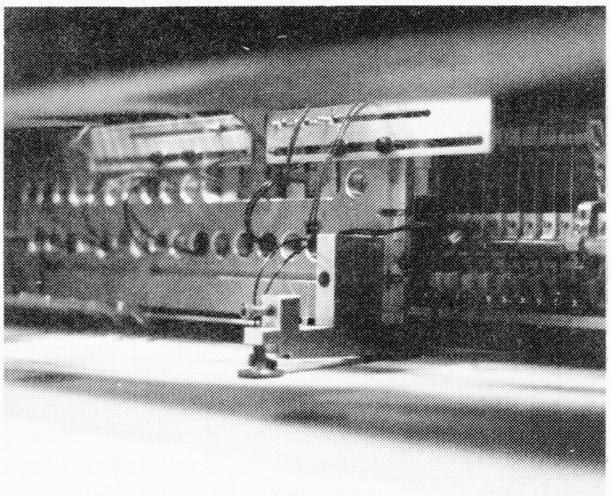

The first section of the action inside the piano during a preliminary test.

specification of minimum moving parts and autonomous control. The forward vertical rail is the mounting for the hammer solenoids which are staggered in two vertical rows. The rear shorter rail holds the damper solenoids in a similar staggered configuration but this time back-to-back on either side of the rail. The front rail has circular holes drilled at intervals to allow the wiring that runs between the solenoids and the driver boards to pass around any of the moving parts and also reduces the weight of the system.

Each of the three sections has at least one locating foot. For the full action length there will be five: one at each end and three located a various points in between. Since each module locks together, two feet per module are unnecessary and only add to the complexity of installation. Feet can be added or subtracted during an installation depending on the particular instrument. They are also locatable anywhere along the mounting rail where the damper levers are absent. This will usually occur when the frame struts separate the strings into sections.

The solenoid cores (the iron part which moves under the influence of the magnetic field) will move 15 millimeters for hammer operation and 5 - 7 millimeters to lift the dampers. The hammer and damper movements are actually in contrary directions, that is, the hammer solenoid cores move up while the damper solenoid cores push down.

The dampers are operated in the same way as the original action. The solenoid pushes down on a lever which raises the damper. It is expected that these solenoids, while on for considerable periods of time, will require less power. Consequently they will not heat up so quickly nor be subject to aberrant electro/mechanical behaviour such as core 'chatter'.

The action was not conceived as being a permanent part of any particular piano. The full implications of this desire for portability emerged as the research progressed and it was soon realised that it was going to take a considerable amount of time to assemble in an instrument. More over, it was going to be tedious and frustrating due to the massive amount of adjustment required to align damper levers and hammer solenoids. However, to make the task easier, much of the adjustment and assembly was envisaged as taking place before the action was installed. Every attempt therefore will be made to alleviate complications with the wiring and final adjustments before it enters the instrument. A clearer understanding of solutions to these types of problems are likely to result from installation experience rather than speculation.

The action frame was machined from 16mm and 25mm aluminium plate. This provides considerable rigidity and yet is relatively light compared to the all up weight with the solenoids.

Testing and Performance Expectations

Once the mechanical component is completed the system will be tested with existing electronic hardware to verify that it is operationally successful. This testing stage will not initially take place in an instrument since that adds unnecessary complications. When the system has satisfactorily completed the initial tests and been examined for any potential problems, it will be installed within a piano and the testing will begin again under the intended operational conditions.

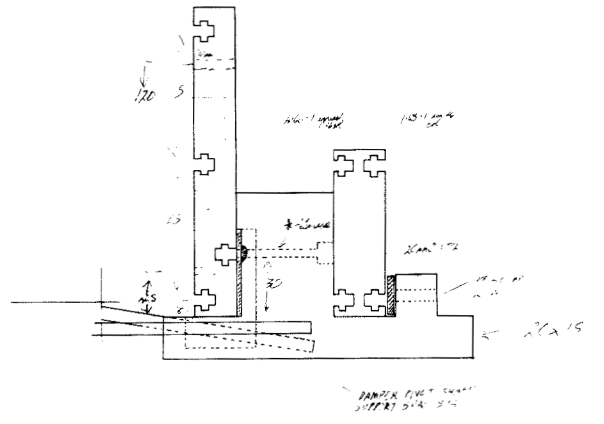

Figure 3. A cross section of the action frame.

Underlying the success of the project will be its ability to perform to fundamental expectations. Since the movement of each part of the action is very small, it should in theory at least, function with considerable rapidity. This has been proven to some extent on an earlier system (Riddell 1988). The particular layout of the solenoids (vertically) introduces no new forces other than the existing functional ones of electro-magnetism and gravity. However, rapid operation may be inhibited due to the sophistication of some of the co-ordinated operations of the mechanism, where the dampers and hammers are able to operate synchronously or autonomously during a performance. In a normal operating mode, for example, a damper should be free of the string by the time the hammer strikes. The system is therefore required to have the damper solenoid activated in time to avoid the hammer striking the string while the damper is still relatively close.

Once the mechanical component has been tested, attention will turn to the most critical aspect the system — real-time control. An initial consideration of real-time operation results in the view that a distributed approach may be necessary. That is, where the management of complex scheduling and event structures is made easier by partitioning and distributing computation and I/O to subordinate processors. This simplifies and solves problems in one area but unfortunately introduces problems in others. Nevertheless, while control of one piano may not cause major problems, a multiple instrument configuration would at least increase the complexity of control by an order of magnitude for a single microcomputer system and thus justify distributed processing.

Irrespective of what approach is adopted, the system will remain subject to change if circumstances warrant it. Consequently, every effort is made to maintain flexibility throughout, hopefully allowing change to one area without necessitating extensive change in others.

Software is particularly suited to modularity and revision. If, for example, the lower level processes manage the note by note activity they need not change when software is changed at a higher level. This is a common practice these days in software development but difficult to carry out if the lower level functions become to embedded or too specific.

Conclusion

From the outset of this project many obstacles and problems arose that made it appear impossible. It was not simply the scope of the project but the perceived inherent difficulties within each major section. Months later when those obstacles can be considered with some hindsight, it is clear that some apparently insurmountable dilemmas were resolved with less effort than initially imagined. Therefore to have even arrived at this point in the project is very encouraging and although much work remains to done, the next stages can be viewed less as a series of problems to be faced and more as milestones to be passed.

Footnotes

1. I gratefully acknowledge the financial assistance of the Australia Council in this project.

2. I would like to thank Professor Keith Cole, Dr Ron Miller and Horst Dressel of the Department of Physics at La Trobe University for their co-operation and encouragement in getting this project underway.

3. The action owes its existence to Dr Micheal Podlesak who contributed his expertise and interest during the experimental stage and Marshall Maclean of the Physics workshop, through whose skill and insight the action takes its present form.

Bibliography

Apel, Willi. **The Notation of Polyphonic Music. 900-1600.** Cambridge, Massachusetts: The Mediaeval Academy of America. Publication No. 38. 1953.

Grover, David S. **The Piano. Its Story from Zither to Grand.** London: Robert Hale. 1976.

Riddell, Alistair. *The Computer Controlled Piano: New Performer, New instrument.* **NMA1**. Melbourne : NMA Publications 1982. pp.6-9.

Riddell, Alistair. M.A. Thesis (in progress). Department of Music. La Trobe University. Bundoora. Victoria. 1988

On not having read *Think — or perish! Towards a confident and productive Australia*, by Donald Horne, Commission for the Future Occasional Paper

Chris Mann

The body, being water, magnifies sound. Speech resonates. Internally it affects organs, externally it affects muscles, acupuncture points and the like. Reverse this. Use a network of surveillance cameras to model a body in space. The model is a map of acupuncture meridians and points, muscle areas, organs, dispositions such that movement of any one or set of these points triggers a speech event. (Speech events being suffixes, prefixes, particles, stress patterns, intonations, conjugations, learning patterns, narratives, whatever.) Now, add another figure and associative logics. Articulate. This is one do-able two cheaply three you've just invented a homunculus.

A bunch of residents tried to set up memory-free learning (they didn't agree to do this, by the way, they just did). Idealism packaged for export (the functionalism of legitimation) is at least a public choice: the loyalty of guilt has the same relative value (utility) as a stool pigeon and knows exactly what it would like to have said to that client of pathos bad questions could call taste. (I dream too that my teeth are falling out. (On speaking with your mouth full: usury is one continuous rhyme.)) A language of negatives (compulsion neurosis) takes itself seriously by pretending to be the uderdog, a password of boredom (and therefore read): a bed. How odd of god to flog the news. Where work makes sense, unionise sense (where capital is organised, strike). Greed: the oops theory — truth in a basket — (it's ok, we forgive) is too whatjamacallit, too.

The advantage of art over life: art is not restricted to good ideas. (In those days it used to be smart; these days it's smart to be sorry.) This would have been perfect (It's not that we have anything to say but that something needs to be said (I define as noise the fact that you haven't said anything, yet (yet of course is poison).) Yes-buts don't grow on trees you know.).

I told you so is just a play on words. What's that? is not.

Composition as a problem-seeking environment;

On explaining to a bureaucrat that they have a problem:
What do you mean independent? Iconoclast? Tradition? Whose? Criteria? What is self-evident about mediocrity? What is more traditional than experiment? How many internal contradictions does it take to screw in a light bulb?
Arts functionaries who are known to support that which they understand: (a do-it-yourself list) Arts functionaries who are known to think: (a do-it-yourself chain letter) (And if the question comes back Well, how would you organise it? Who's asking? Why? Do you want a job?) Meaning is contaminated by theory, by what is held to be true (truth is a propositional dictionary (look up aesthetics (relevance spells charm with an R))). And meaning is organised, it is one of those organisations of I-told-you-so where utterance cannot be subject matter. I, for one, will make sure she is spoken to about it.

On alchemy as a rubber stamp (no thing makes sentences true):

prediction is a given/ accounts for in a fit/ coping is here evidence/ available in pink
Therefore meaning is a charity of punctuation.
The resentful hunch of a peevish memory (an explanation of criteria) that says misunderstanding is no opinionated husk of commerce:
If Descartes writes in French for women and Hegel teaches philosophy how to speak German, Wittgenstein is corrected not by Chomsky but by speech (a sentimental bait (the negative of sowhat)) and doubles up by pun a competent subject of propositions defining romance as a greed that's so far up itself it's stuffed. (The transcendental bootstrap of profit is beauty in a narrative of there and then. A blushing corpse is a qualified conversation piece.)
Accounting (suburban surveillance — if you are in the army doesn't that count as unemployed?) is an ideal that wants to send negation to coventry for breaking step and looking sideways in a lift at someone's lunch (an object is the opposite of shock(ed)). Management is a machine to generate perspective (a pretence) of a sort (case endings of the rich and famous) that read. Please.
The charm of error: (bothering the false) a conceit (self confidence of image) that does attract, a flak. To report (endorse) the mimic (the notation of the novel) as indifferent denies the paranoid due dull, a bloat. Oh, our (noblesse oblige in a soap) frog in the throat that impersonates the automatic by falling on a loan, show up. An exaggeration that knows itself (effect) is dead (affect) to tends: it rather infects ends.
a half truth: (learning) (another) language
a neurosis (where a lever is a line of flight (opening (criticism by dimension))): a lot of
an alibi (technology as boutique accident): oughta there is no irony in doubt: service delivery (war) is hostage to the prosaic punchline of taste
keep it thin: locality (stodge) is just a oncer (easy pickings (inertia — a thing is just an economic lie)) and the reliability of the military consists in that we just don't know (a model of nowhere)
terms (the bits of propositions (ethics)): (it's a on the tip of my tongue) pun a dictionary was particularly not license to not know: blushes at the very structure of thought (to all intents and purposes a possible), a subtle agreement of existence (read: intent)
to desert (sit down, — the assassin has no standing orders) (a lunch box on teaching dillettantes how to fart): a rip-off of a wall (but no police came coz they didn have police then)
chance (the shadow of map) and the blitzkrieg of reason: over there (an agent of being seen to be on a false errand (attack)) where information is training by deskilling (learning as behaviour, some admissable will)
yes, I understand: yes, I agree
gesture is that opinion makes a comfort into science (the rule of law)/ and so frustrates the metaphor of defeat (logistics by belief)/ a sub clause that mimes so there (the anxious general) as fear, is sweet:/ money sells information, the military sells doubt, wait for the aristocracy to package silence
non linear dynamics goes to town (two bob each way): a shaggy dog story; take a letter, turn it through three hundred and sixty degrees, portion it out as an alphabet and to avoid the compromise of words (junk bonds) neglect the vowels. Hebrew as the classic number cruncher. Zero is just a philosophical asprin out of Sanskrit and Latin some controlled extrapolation, a knock knock joke for (electronic) communion.

PS

now that I am a priori (and got a right to want):/ tee is for tautology and barracks at the front/ aich is for a ransom all done up like a do/ ee is just a license fee that subs for flu/ and tee aich ee spells you
to intuit (pleasure (begs the question) bags the diff) the abstract as distraction to the power of vindictive is here sufficient if (a known) that supposes con to be a noun (now now) of this: too much monkey business
that we agree is premissed on agreement: too bad looks good: non-contradiction as a kind of possibility (essentially)
because 'that' is a doing word: because
to be for: a case of thens
it: so
method as the means of en-maybe-ing a curly predicate (a belong (his ism is hymns)): a zom-bie, or
a most of: nothing much
and there are those metaphysical punks that hold all givens to be originals (that representation is some sort of bridge (a sentimental point) or bargain) that technique. It renders memory to compare a whiff of where metaphors are dressed as wise, a up-the-ante set that might comprise a foot. Nouns are functional apologies (on a bond) and beg to differ (an obligeing pragmatism (on the blink)): a self-reference is only an assumption à la charm, a its. Busy is 'e? Choice (a toy, a failure of nerve) really please. Aye aye. (Crap of course is a logical tunnel, a cash. (Descartes was a failed mercenary, Foo was right.) From which we deduce what we lack. (Or law and order.)) Now is not very long (just a proof looking for the exit), a fashion (pardon) or sentimental irony, a blush. Repetition

only ever gets to be reasonable. Send more money.
In that simple sentences, statements are taken to be propositions the self-reference of information gives us science (formalism, a knowing smile) (credit). It defines the actual (a contingency, a necessary and sufficient future) as a radical ignorance (negation is only two dimensional otherness, a parochial) understood. Certainly, I mean, there is still room — one potato two potato three potato four places at table and you mash them. A normative ply (doubt (a paranoid assertion)) is just that. Complacent is a form of analysis. The manners of not. (And on the romance of counting: makes up.) Infatuation jails that distance we'd rather sell. Frankly (with intent) imaged is as this, some good ghost. (Post.) Damn, I forget. A droll promise (information is just cheap PR, a wrap (a the ad) (and on the Spruce Goose and other paper planes (no radar print): a rayon shirt is always a good buy). The vanity of sympathy (data) is so disposed to fickle fuck we see no such accrual as might sustain a mouthing off (a cliche is not a douche (language is only a metaphor for size)): after all (suits me) a gather gathers (and even then with dicky teeth). Darwin was trying to prove etymology. A cinderella echo. See Saw and Shutup went to the sea. A seditious compromise. Talk is cheap. Copy is free. Use is payment.
when you come to think of it one is pretty funny (a crusading positivism of puns): ha ha ho ho he he yes that's it he's the one blame him (original sin)
a couple of grammars went bowling and this is what they said, come in spinner come in, a fence is a trader a trader is a wall walls have ears and ears are dead but fences
fill me in on fiction (the dapper fag) and fuck off: from subject to verb without a likeness, lippy the diction (a flapping dag) and fuck off
oh very well (grudge), anyway sounds ok. This and that met something. They discussed doing stuff. And didn't. When the time came however it was on. Too bad. But not too bad. Adjectives are at least learnable by small people. Pretty much. must just (know-all) was name calling and wanted to talk down one two sure (don't answer back) and the agent (mistaken for a had-to) laze do may (enough): a doodle is a noodle on a lead (nice one) pad pat
Does it matter (bad) that cause is merely flag and on (good) condition (anglo) that probly sue the premise for its should and bangle? Marches? Blameless lapses? How can you tell? Lets please explain. Might've was quite right (at the time) (need likes to think that it's somewhere else) and bets on mine. A stitch. Out of it. (I know.) Contradiction is a one dimensional proof (the best of all possible rules (juve) deny (so) 'as' to be consistent) of some. False (defined as price) is only as intitial condition of context (the art of partly) and next to what the logical call properties. It holds. A they is a also-ran of cans. Still looks the same. A predicate. (The profit of statements.) And better.

Exercise 1: For for example five players, say every fifth word (in turn) so as to maintain (prosodic) sense. Fluently.
Exercise 2: Establish other (than circular) patterns.
Exercise 3: With an ear to timbral and intonation qualities allocate voices more particularly (undo excercises 1 & 2).
In performance use throat (contact) mikes and (sealing) headphones so that each performer can hear the others. However no performer should be able to hear themselves. The mike should trigger a delay such that when a performer voices she only hears what was just said (by the previous speaker).
if and then are descriptions what seem to be the case (although it does imply the bribe) a qualitative paste unlesses could and puts it back in place (reference but)

BEYOND THE BODY:
AMPLIFIED BODY, LASER EYES & THIRD HAND.
STELARC- ARTIST: YOKOHAMA INTERNATIONAL SCHOOL

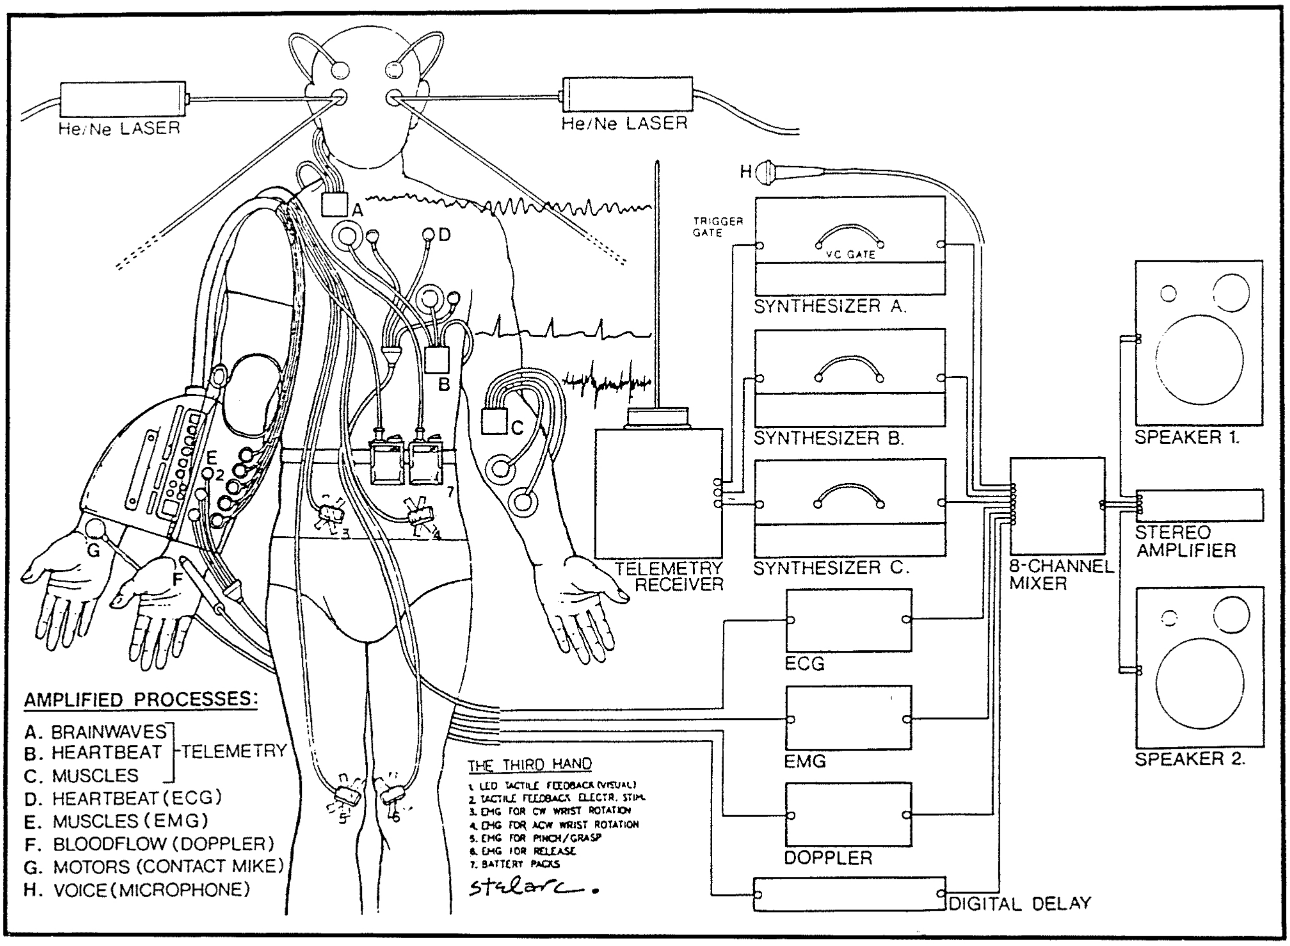

THE INVASION OF TECHNOLOGY: MINIATURIZED AND BIOCOMPATIBLE, TECHNOLOGY IMPLODES BACK TO THE BODY, NOT ONLY LANDING ON THE SKIN BUT EMBEDDING ITSELF AS AN INTERNAL COMPONENT. IMPLANTED TECHNOLOGY ENERGIZES THE BODY, ACCELERATING IT TO ATTAIN PLANETARY ESCAPE VELOCITY. EVOLUTION ENDS WHEN TECHNOLOGY INVADES THE BODY. IT IS NO LONGER OF ANY ADVANTAGE TO EITHER REMAIN "HUMAN" OR TO EVOLVE AS A SPECIES. HUMAN THOUGHT RECEDES INTO THE HUMAN PAST. THE END OF PHILOSOPHY, THE END OF THE HUMAN FORM.

1. If the earlier events can be characterized as PROBING and PIERCING the body (the three films of the inside of the stomach, lungs and colon/the 25 suspensions), then the recent performances EXTEND and ENHANCE it. The amplified internal rhythms, laser eyes and mechanical hand acoustically and visually expand the body's parameters. They can no longer be seen as biofeedback situations (they never really were) but rather SCI-FI SCENARIOS for human-machine symbiosis -with sound as the medium that reshapes the human body, for redesigning an obsolete body. It may not yet be possible to physiologically modify the body, but it can resonate with modulated rhythms. The body does not simply aquire an acoustical aura -its humanoid form is stretched and restructured with sound. The amplified body is no longer the container of its rhythms. The humanoid form is transformed into the cuboid space. The body becomes hollow, resonating with its own echoes.

HOLLOW BODY: OFF THE PLANET, THE BODY'S COMPLEXITY, SOFTNESS AND WETNESS WOULD BE DIFFICULT TO SUSTAIN. THE STRATEGY SHOULD BE TO HOLLOW, HARDEN AND DEHYDRATE THE BODY. EXTRATERRESTRIAL ENVIRONMENTS AMPLIFY THE BODY'S OBSOLESCENCE, INTENSIFYING THE PRESSURES FOR ITS MODIFICATION. THE SOLUTION TO RADICALY REDESIGNING THE BODY LIES NOT WITH ITS INTERNAL STRUCTURE BUT WITH A CHANGE OF SKIN.

2. The artificial hand, attached to the right arm as a third hand, is capable of independent motion, being triggered by the EMG signals from the abdominal and thigh

muscles. It has pinch-release, grasp-release, 290° wrist rotation (CW and CCW) and a tactile feedback system for a "sense of touch." But whilst the body activates its extra manipulater, the real left arm is REMOTE-CONTROLLED -jerked into action by a muscle stimulater with varying intensity of voltage and rate of frequency, both in random and repetitive modes. Of necessity, this remote-controlling is done intermittently (it is quite painful) and is used to "pace" the body's performance and to alter the body's general condition, thereby affecting its acoustical field. The stimulater signal is used as a sound source, whilst the motor mechanism of the Third Hand is picked up by a contact microphone.

ANAESTHETIZED BODY: THE BODY INSERTED INTO THE MOBILE MANIPULATER UNIT WILL SPIN, GLIDE, CIRCLE AND HOVER. ITS MECHANICAL ARMS WILL BE OF PRIMATE PROPORTIONS, DOUBLE-JOINTED AND CAPABLE OF HIGH-SPEED MODES OF OPERATION. AS WELL AS EMG CONTROL IT WILL ALSO HAVE AUTOMATIC COMPUTER CONTROL AND A SOUND ACTIVATION INTERFACE. IT WILL NOT SIMPLY AUGMENT BUT RATHER REPLACE THE HUMAN LIMBS. THE BODY PLUGGED INTO MACHINE SYSTEMS NEEDS TO BE PASSIFIED. IN FACT, TO FUNCTION IN THE FUTURE AND TO TRULY ACHIEVE A HYBRID SYMBIOSIS, THE BODY WILL NEED TO BE INCREASINGLY ANAESTHETIZED.

3. Body processes amplified include brainwaves (EEG), muscles (EMG), heartbeat (ECG), pulse (PLETHYSMOGRAM-finger clip-on photo-electric type) and bloodflow (DOPPLER FLOW METER), with a KINETO-ANGLE TRANSDUCER transforming bending motion into a sequence of sounds. A C-DUCER has also been used over the larynx to pick up vibration in the throat and stomach activity has been monitored by swallowing a transmitter (tethered so that it can be later extracted). With the heart, the opening and closing of the valves, the gurgling of the blood and the gushing of the blood thru the wrist can be amplified best by the Doppler ultrasonic sound transducers -the pencil-type probe for deep monitoring and the flat-type for the shallow wrist section. Although the pencil-type probe has several disadvantages in having to be held and needing intermittent application of gel (over the length of the performance), quite dramatic changes of sound occur over a small change of skin scanned and by pointing the probe at slightly different angles. By constricting the radial artery of the wrist, the sound varies from the normal repetitive "whooshing" to a "clicking" as the blood is dammed with a flooding rush of sound as the wrist is relaxed. The use of a TELEMETRY UNIT minimizes the hand-wiring of the body (transmission distance is 30m) to the

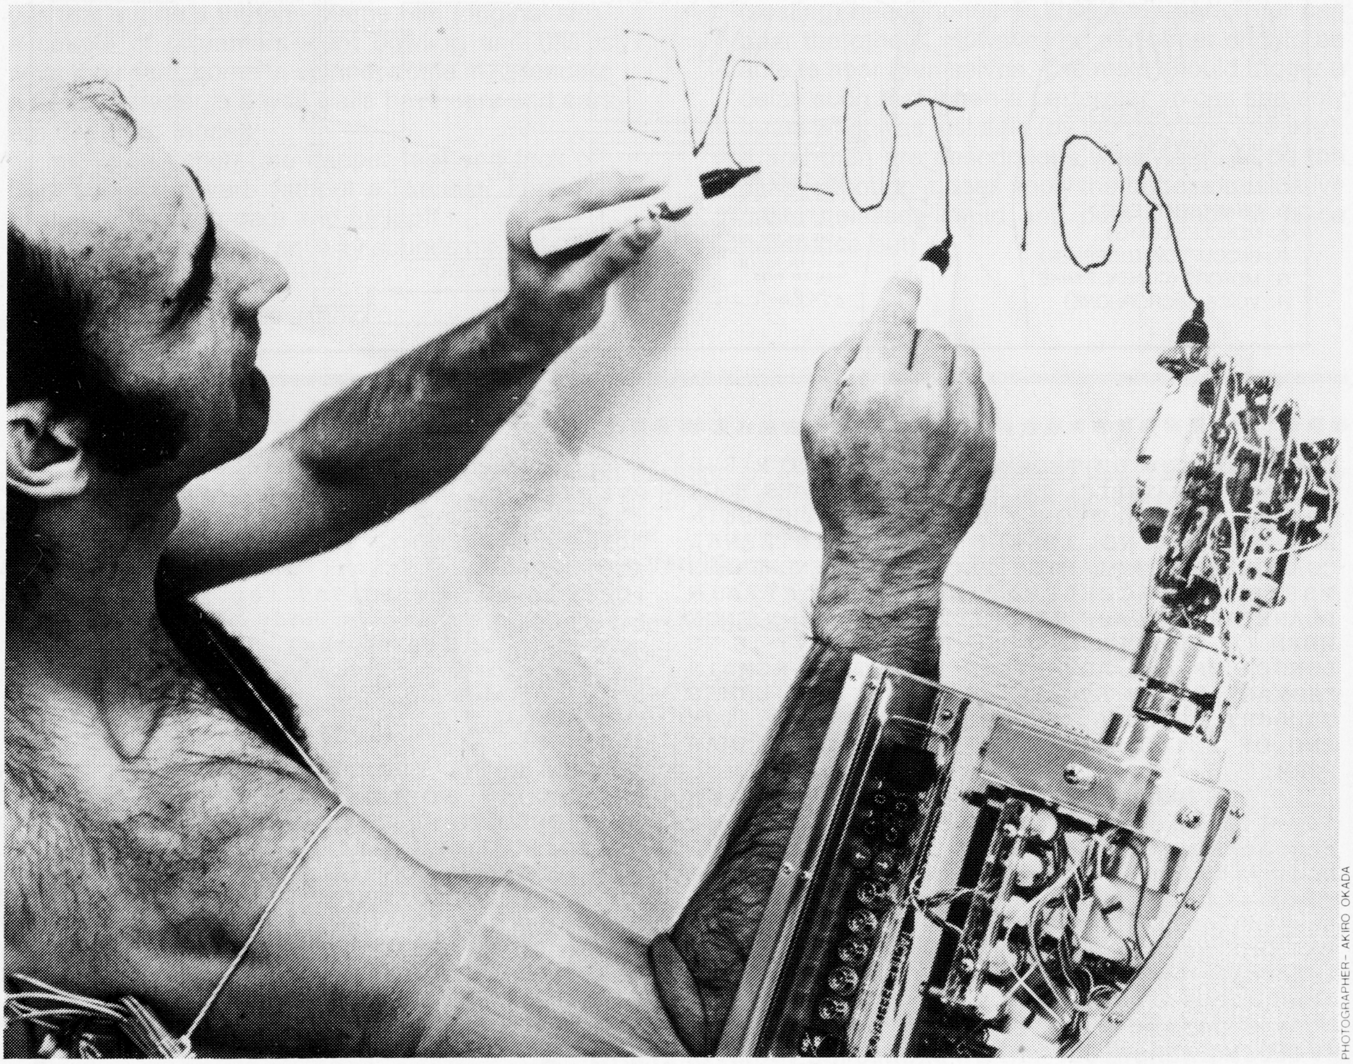

PHOTOGRAPHER - AKIRO OKADA

Stelarc, *Handswriting: Event for Three Hands* (1982)
Australia: Nine Contemporary Artists

June 30 - August 14, 1984, Los Angeles Institute of Contemporary Art
An Official Program of the Olympic Arts Festival, Los Angeles, 1984

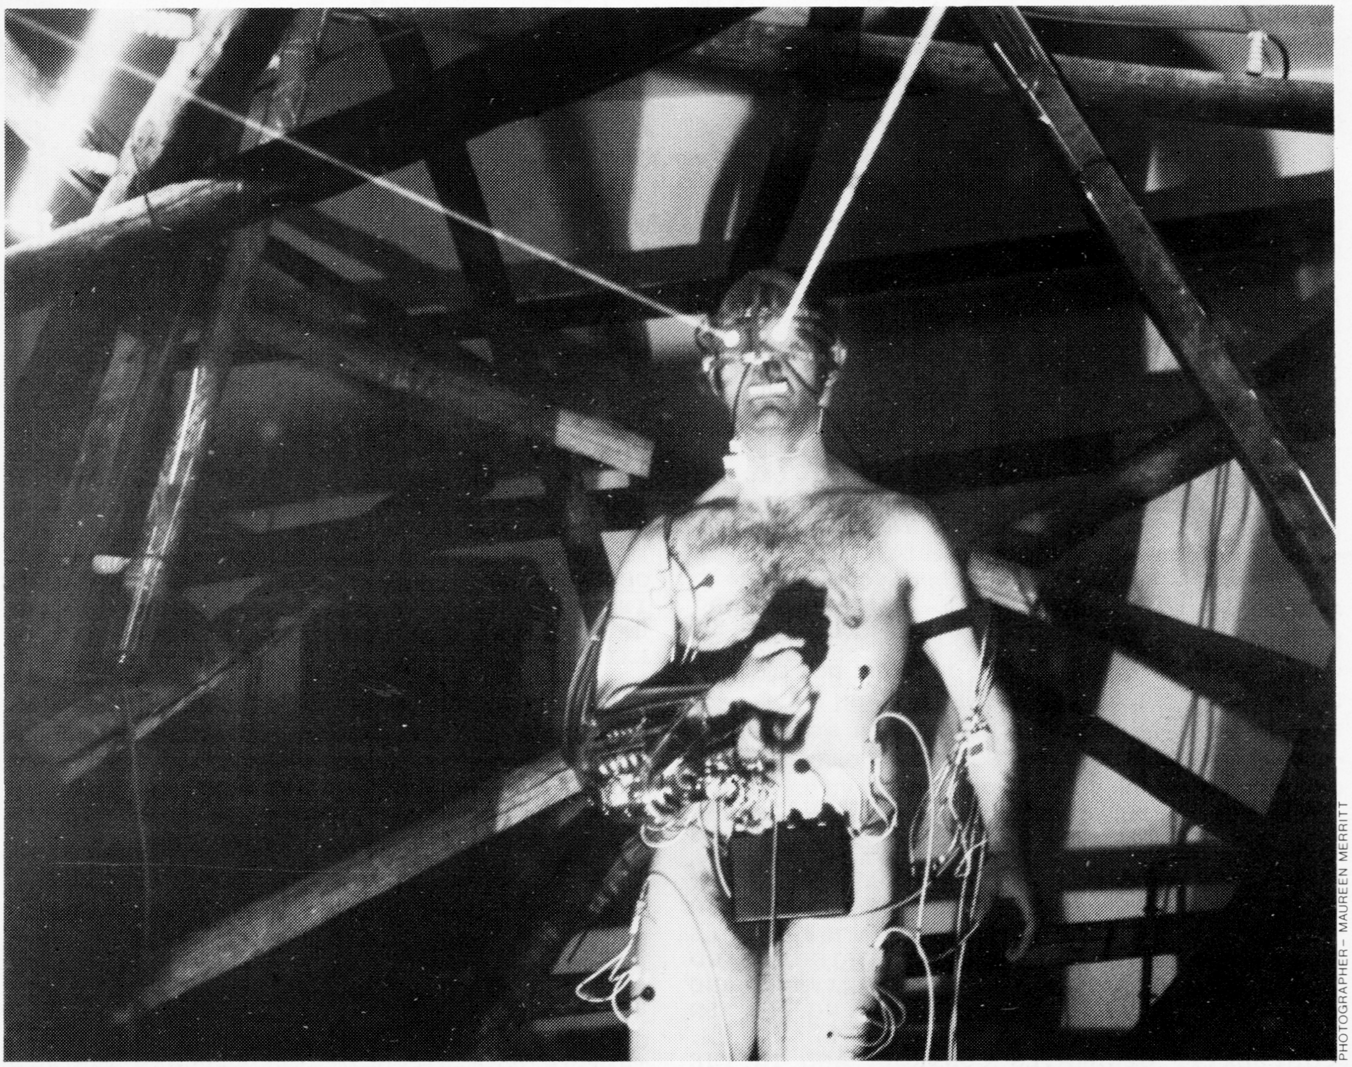

EVENT FOR AMPLIFIED BODY/ LASER EYES AND THIRD HAND MAKI GALLERY, TOKYO- 2 MARCH, 1986.

equipment, safely isolating it from the electrical system, removing possible hum and noise and allowing the body freedom of movement.

HUMAN-MACHINE SYMBIOSIS: REMOTE SYSTEMS AND SURROGATE ROBOTS PRESENT THE GREATEST POTENTIAL AND THE MOST INTRIGUING DILEMMA. TELECHIRIC SYSTEMS WOULD HAVE TO BE MORE THAN HAND-EYE MECHANISMS. THEY WOULD HAVE TO CREATE A KINESTHETIC SENSE IE. PROVIDE THE SENSATION OF POSITION, MOVEMENT AND BODY TENSION. THIS PRESUPPOSES SOPHISTICATED HUMAN-LIKE AND INTELLIGENT ROBOTS CAPABLE OF SOME AUTONOMY EVEN WITH HUMAN PARTICIPATION IN THE LOOP. THE PROBLEM IS WHETHER SURROGATE ROBOTS CAN ADEQUATELY SENSE AND ACT -COLLAPSING THE TIME-SPACE BETWEEN THE BODY AND WHAT IS PERCEIVED AT A DISTANCE.

4. Actions such as flicking of the fingers, bending an arm, twitching the facial muscles, turning the torso and lifting the leg bring forth a cascade of sound. Powerful acoustical effects can be generated both by discernible gesture and invisible internal contractions and control. The sound field is configured by buzzing, warbling, clicking, thumping, beeping and whooshing sounds. A combination of percussive-like and wind-like sounds; of triggered, random, repetitive and rhythmic sound. There is a general score or structure in the performance depending on the number and type of body frequencies amplified. Within these performance parameters the body improvizes depending on the feedback it generates. Orchestration of the event involves selective tuning into/out of channels of sound (varying the complexity); increasing or decreasing the volume of certain sounds (contouring the sound field); physical control of certain body functions and motions; activation of the mechanical hand and the use of digital delay (foot pedal) to loop and superimpose sequences of sound. The general dilemma of the process is to modulate the original signal in a way that best reflects the body function and mantains an identity with it. An interplay between physiological control and electronic manipulation.

THE HUM OF THE HYBRID (NO BIRTH/NO DEATH): DEATH DOES NOT "AUTHENTICATE" EXISTENCE. DEATH IS AN OUTMODED STRATEGY REQUIRED OF AN EVOLVING SPECIES. IT IS OF NO ADVANTAGE TO THE AWARE INDIVIDUAL! TECHNOLOGY EQUALIZES THE PHYSICAL POTENTIAL OF HUMAN BODIES AND STANDARDIZES HUMAN SEXUALITY. WITH THE POSSIBILITY OF NURTURING THE FETUS OUTSIDE THE WOMB THERE TECHNICALLY WILL BE NO BIRTH. AND IF THE REPLACEMENT OF MALFUNCTIONING PARTS CAN BE FACILITATED THEN THERE WOULD BE NO REASON FOR DEATH.

THE MODIFIED BODY WILL BE ASEXUAL AND IMMORTAL. THIS IS NO MERE FAUSTIAN DESIRE NOR SHOULD THERE BE ANY FRANKENSTEINIAN FEAR. REDESIGNING OUR BODY MEANS REDEFINING OUR ROLE.

5. In previous events, He/Ne (Helium-Neon) lasers were reflected off small optical mirrors stuck to the eyes. This was simple needing no other paraphenalia but it required the head to be almost totally rigid and always facing in one direction. It limited the laser sequences to short durations and to only a direct frontal effect. Now Ar (Argon) beams are propergated thru OPTICAL FIBRE CABLE and an input-output lens system -allowing more powerful lasers to be used safely, with the head and body being able to turn without losing the beams. The output lens are positioned in front of the eyes by an aluminium-frame head structure to allow the eyes to track the beams. The laser eyes are modulated by the heartbeat, pulsing on and off -the sounds of solenoid clicks amplified to synchronize with the ECG. By blinking, twitching facial muscles and oscillating the head it is possible to SCAN the space and SCRIBBLE images, seemingly with the eyes.

DETACHED BREATH/SPINNING RETINA: **THE PSYCHO-SOCIAL FLOWERING OF THE HUMAN SPECIES HAS WITHERED. WE ARE IN THE TWILIGHT OF OUR CEREBRAL FANTASIES. HUMAN EXISTENCE CAN NO LONGER BE JUSTIFIED "IN ITSELF." THE TECHNOLOGICAL TERRAIN CONCEALS COUNTLESS BODY PACEMAKERS -VISUAL AND ACOUSTICAL CUES TO ALERT, ACTIVATE AND CONDITION THE BODY -DIRECTING IT IN PRESCRIBED DIRECTIONS AND VELOCITIES. COMPLEXITY GENERATES CONTROL. AND IN THE HIGH STIMULATION OF CONTEMPORARY SOCIETY THE REFLECTIVE MOMENT BETWEEN INTENTION AND ACTION IS ERASED. THE SINGULARITY OF COMPLEXITY DISINTEGRATES THE PERIMETER OF COHERENCE.**

6.- The installation, often of large rocks, suspended poles of wood and tensegrity constructions manifests mass, weight and gravity emphasizing the physicality of the body -providing the setting for its acoustical transformation. The installation is activated when the body is plugged into it. The body performs in a structured light environment, which flares and flickers, responding and reacting to the electrical discharges of the body -sometimes synchronizing, sometimes counterpointing. Light manifests and further amplifies the body´s internal rhythms. It does not simply respond and illuminate but is understood as a physical phenomenon that can it turn directly affect certain body rhythms. For example strobe flicker triggering and driving brainwaves. The light installation not only extends the body but also helps to redefine its form.

OBSOLETE SKIN: **SKIN HAS BECOME INADEQUATE IN INTERFACING WITH REALITY. TECHNOLOGY HAS BECOME THE BODY´S NEW MEMBRANE OF EXISTENCE.**

7. In amplifying the body, the audience is immersed in its rhythms. The body does not merely acquire an acoustical aura but rather the audience finds itself inside the body. By externalizing internal rhythms the distance between performer and audience collapses. It is a post-verbal communication where the audience identifies instantaneously with the synchronized sounds from posture and gesture.

TOWARDS HIGH-FIDELITY ILLUSION: **THE SIGNIFICANCE OF TECHNOLOGY MAY BE THAT IT CULMINATES IN AN ALIEN CONSCIOUSNESS -ONE THAT IS POST-HISTORIC, TRANS-HUMAN AND EVEN EXTRATERRESTRIAL.** □

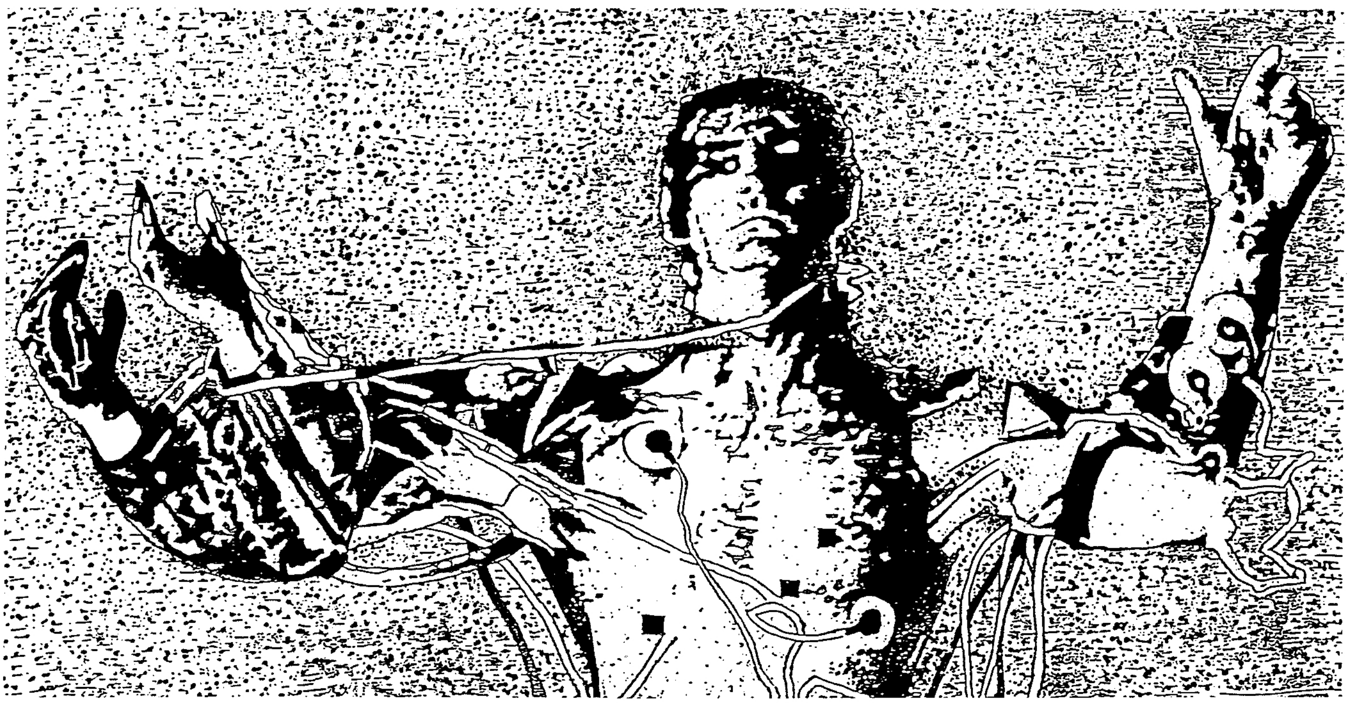

Monophonic Variations

Spontaneous vs pre-meditated composition

Greg Schiemer

Monophonic Variations is a work created under the co-authorship of myself and Graeme Leak. In this work, a solo performer and a computer play interactively improvising on a limited set of rhythmic ostinati. The computer generates a repetitive, monophonic sequence of notes, produced with varying degrees of repetition and at different speeds. The performer can either interact with this process to influence the way that the computer generates the notes, or improvise a counterpoint to the computer-generated note line.

The work was realised using a low-cost computer kit called the DATUM, manufactured by Gammatron in South Australia, and which I modified in 1984, making changes to the circuitry and adding hardware to enable communication with MIDI equipment[1].

The modified DATUM makes an extremely versatile and low-cost machine for the musician who is a novice to machine language programming and who wishes to redefine with musical vision and imagination the constraints under which compositions can be written. It has allowed me to find what I am able to do with MIDI rather than follow the musical agenda set by Japanese and American corporations. It has given me 'bush-tucker' I'd never find at Macdonalds.

Monophonic Variations was programmed — in 6802 machine code to function in less than 350 bytes of ROM — to communicate with a MIDI percussion set-up devised by Graeme Leak, which includes a Roland Octapad and an AKAI sampler. Graeme, likewise has not confined himself to the role of performer. He too has challenged the hidden compositional aesthetic implied by the designers of MIDI equipment through his handiwork in carpentry and soldering, and devised additional playing surfaces in order to extend the capabilities of the Octapad, as well as collecting and sampling all the sound material for the piece.

In *Monophonic Variations* the performer has no score. In one sense the organisation of the work is intuitive, relying largely on the performer's creative musicality for its shape and continuity. However, there was pre-planning of the work which, in retrospect, falls into three areas. These are :

- the computer process that generates the notes
- the choice of how the performer affects this process
- the configuration of the devices in the system to allow the AKAI Sampler to be played by both performer and computer.

The computer in *Monophonic Variations* has no itinerary of notes to be played, literally producing notes on the spur of the moment in performance using a process called a Polynomial Generator. In this way, not only are the notes un-prepared — in the sense that notes are prepared for a performance by writing a score — but the process can be altered at the whim of the performer to produce notes in a different fashion.

How the computer produces notes, and how the performer affects this process is described below

This article will describe the binary processes used in *Monophonic Variations* prior to discussing what I mean by Interactive MIDI. As all digital processes being described are operations performed on binary numbers, it is therefore best to describe these using a short-hand form of binary, Hexadecimal. MIDI Note Numbers are offered in decimal only at the final stage, for easier cross-reference with the staff notation in the illustrations.

Appendix 1 is a cross-reference between Hex and Binary.

Generating Notes

The Polynomial generator was written for the DATUM to simulate the operation of a hardware sequencer which I built in 1982 for controlling the Tupperware Gamelan[3].

The sequencing process implemented on the Tupperware Gamelan uses a handful of operations performed iteratively on a pattern of eight binary digits, to produce a variety of elaborate hockets. In *Monophonic Variations*, 6802 machine code instructions were used to simulate these operations and generate polynomial values which were then used by the computer to play notes.

The process for generating these note values — rotating a bit pattern to produce a sequence of MIDI notes, varying the feed-back to alter both the melodic contour and the sequence length, and filtering bit patterns in a sequence to alter the melodic contour — is described below.

Rotation and Masking

Rotating a pattern of bits can produce a repetitive sequence of notes. It is an operation that is performed on the contents of a digital device called a Shift Register. Such a device typically consists of

eight cells, each capable of storing either a '1' or a '0'. Rotation causes the contents of each cell to move sideways to the adjacent cell, with the bit shifted out at one end connected back into the register at the other end.

Whenever rotation takes place in a microprocessor, the eight-bit register plus the Carry bit re-circulates like a nine-bit register and a repetitive sequence is produced, as shown in Figure 1.

It is a simple matter to convert this rotating sequence into notes. Ignoring the Carry bit and the most significant bit of the 8-bit pattern, produces a MIDI Note Number. Like all MIDI Data Bytes, A MIDI Note Number is transmitted as a seven-bit number preceded by a '0'.

Masking every 8-bit number with 7F hexadecimal (0111 1111), using a logical AND instruction, ensures that the most significant digit is always a '0'. For example, if all the 8-bit patterns in Figure 1 were masked with 7F hex, they would produce a cyclic pattern of MIDI Note Numbers shown in Figure 2[4].

step	8-bit register (un-masked)	Carry	8-bit register (masked with 7F)	MIDI Note Number Hex	Dec
1	1000 0000	1	0000 0000	00	0
2	0100 0000	0	0100 0000	40	64
3	0010 0000	0	0010 0000	20	32
4	0001 0000	0	0001 0000	10	16
5	0000 1000	0	0000 1000	08	8
6	0000 0100	0	0000 0100	04	4
7	0000 0010	0	0000 0010	02	2
8	0000 0001	0	0000 0001	01	1
9	0000 0000	1	0000 0000	00	0
etc	1000 0000	0	0000 0000	00	0
etc	0100 0000	0	0100 0000	40	64

Figure 1 *Figure 2*

Varying the Contour and the Sequence Length

It becomes possible to vary a sequence of bit patterns if the Carry bit fed back to the input during a rotation is not a literal copy of one single output, i.e. the least significant (right end) bit, but turned into a '1' or a '0' by a feedback process that represents the combined condition of some of the other bits.In Figure 3. a single feed-back bit is generated from the outputs of bits 6, 5, 3 and 2 using four Exclusive OR gates. When it is shifted right on the next step, a new bit pattern is formed[5].

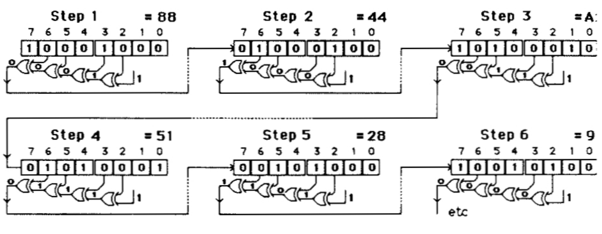

Figure 3

This feedback process is illustrated in 9 stages, though it takes 63 stages before the pattern repeats. Figure 4 shows all 63 Carry bits and 8-bit values generated together with the filtered MIDI Note Number both in hex and decimal.

This is the process *Monophonic Variations* uses to produce the cyclic sequence of 63 notes, as shown in Figure 5. Lines above and below the stave indicate the highest and lowest A available on a piano keyboard.

Step	Carry	REGISTER	MIDI Note Num. Hex	Dec	Step	Carry	REGISTER	MIDI Note Num. Hex	Dec
1	1	C7	6B	107	33	0	21	25	37
2	1	E3	07	7	34	0	10	34	52
3	1	F1	15	21	35	1	88	2C	44
4	1	F8	1C	28	36	0	44	28	40
5	0	7C	40	64	37	1	A2	46	70
6	1	BE	62	98	38	0	51	35	53
7	0	5F	43	67	39	0	28	2C	44
8	0	2F	33	51	40	1	94	38	56
9	0	17	3B	59	41	0	4A	2E	46
10	0	0B	2F	47	42	1	A5	49	73
11	0	05	29	41	43	1	D2	76	118
12	0	02	26	38	44	0	69	2D	45
13	1	81	25	37	45	0	34	38	56
14	1	C0	64	100	46	1	9A	3E	62
15	0	60	24	36	47	0	4D	31	49
16	1	B0	54	84	48	0	26	2A	42
17	0	58	3C	60	49	1	93	37	55
18	1	AC	50	80	50	1	C9	6D	109
19	0	56	3A	58	51	1	E4	08	8
20	1	AB	4F	79	52	0	72	36	54
21	1	D5	79	121	53	1	B9	5D	93
22	1	EA	0E	14	54	1	DC	00	0
23	0	75	39	57	55	0	6E	32	50
24	0	3A	3E	62	56	1	B7	5B	91
25	1	9D	41	65	57	1	DB	7F	127
26	1	CE	72	114	58	1	ED	11	17
27	0	67	2B	43	59	1	F6	1A	26
28	0	33	37	55	60	0	7B	3F	63
29	0	19	3D	61	61	0	3D	41	65
30	0	0C	30	48	62	0	1E	42	66
31	1	86	2A	42	63	1	8F	33	51
32	0	43	27	39	etc				

FB=6C

Figure 4

In *Monophonic Variations*, this same process of selecting and combining several outputs of a Shift Register to produce a single feedback bit, can be reduced to a single-byte variable called Feedback, referred to henceforth as FB. In Figure 3. where the bits that are fed back — represented by a 1 in that bit position — are bit 6, 5, 3 and 2, the value of FB is 01101100 or 6C Hex.

Four values of FB are used to produce cyclic sequences, which become the ostinati in *Monophonic Variations*. These are (in Hex):

6C — (sixty-three note; Figures 3, 4 and 5)
90 — (fifteen note; Appendix 2, Figures 7, 8 and 9)
A0 — (seven note; Appendix 2, Figures 1, 2 and 3) and
50 — (six and three note; Appendix 2, Figures 4, 5 and 6)

The cycle length of the sequence is usually associated with the value of FB. However, the same value of FB can sometimes produce several different melodic contours. A melodic contour can

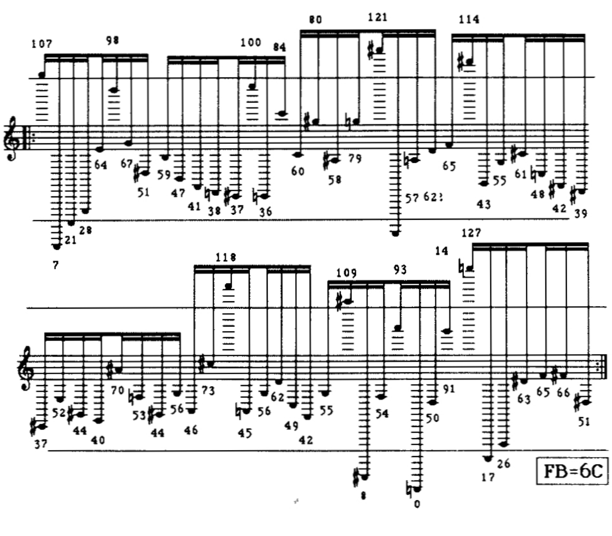

Figure 5

be selected prior to starting the process by initially placing a seed-value in the Shift Register, henceforth referred to as SR. The seed-value is typically any value which occurs on a melodic contour between double bar lines in Appendix 2. The performer 'seeds' SR every time a MIDI Note Number is sent on MIDI Channel 1. If the seed-value is one that belongs to the set of values normally produced by such a feedback process, the function is reset to begin from that point in the cycle. This is analogous to a performer being able to restart a tape loop from any predictable point in the loop — except a percussionist can now use a stick to do it.

If however, a seed-value placed in SR is foreign to the set of values normally produced by such a feedback process, something remarkable happens. The process goes through a transient state where it gradually reduces the 'deviant' values till the point where these give way to 'normal' values, the 'normal' values being those which produce a periodic sequence.

This is almost a model of the way the human memory works, as described in *Mechanisms of Mind* by Edward de Bono, where the human memory process tends to change new information to make it fit the 'moulds' that were created for the retention of previous information. It is as though new musical ideas that arise from the process of improvising, have to contend with musical memory.

Appendix 2 shows examples with a transient sequence, preceding each periodic sequence. The number of steps in the transient sequence is a function both of FB and the amount the seed-value deviates from the values within the cyclic pattern produced by FB[6].

The Polynomial Generator is described in Appendix 3. MIDI Note Numbers are derived from these Polynomial values, by using a MIDI Filter program, described below.

In performance, the Polynomial functions will rarely occur in the pristine form in which they are shown in this article, because of the constant interaction between performer and the system.

Because the value of SR is constantly 'seeded' with each MIDI Note Number produced by the performer, the computer process will tend therefore to deviate momentarily from its cyclic behaviour the more the performer reacts to the rhythms and phrases which it creates. The performer at all times has the potential to introduce that which is foreign to the process, which the process will continually try to assimilate. Variation comes from the tension this creates.

MIDI Filters

Masking the values produced by the Polynomial generator changes eight-bit numbers into seven bit MIDI Note Numbers, but this is only one function of the MIDI Filter. The other function is to alter the distribution of these MIDI Note Numbers. In *Monophonic Variations* these numbers are concentrated in the range between C three octaves below middle C (24 Hex) and E above middle C (44 Hex), though the filter will produce occasional deviations outside this range. This process for restricting the range of MIDI Note Numbers was chosen to get the desired resonance from sounds produced by the AKAI. The filter is really two filters superimposed. One is a 'band-pass' filter which concentrates all MIDI Note Numbers between 24 Hex and 44 Hex. The other filter produces notes that fall outside this range. It is a 'low-pass' filter where the cut-off point is set just above the highest MIDI Note Number, namely 7F Hex. The MIDI Note Number Filter is described in Appendix 4.

Configuration of Monophonic Variations

In the configuration shown in Figure 6 MIDI data created by the percussionist playing the Octapad is used to control the DATUM as well as play sounds on the AKAI Sampler.

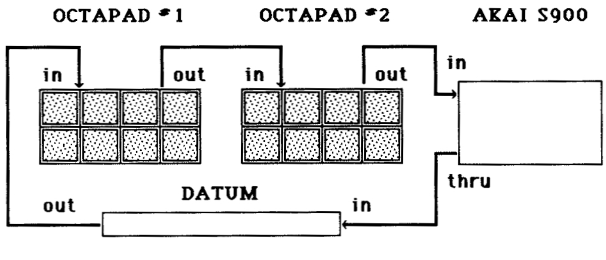

Figure 6

This configuration enables three exchanges of MIDI information to take place:
(1) The percussionist plays either or both Octapads which send MIDI signals that play the AKAI S900.
(2) MIDI signals produced by the DATUM, as well as those produced by the percussionist, play the AKAI S900.
(3) MIDI signals produced by the percussionist, play the AKAI S900, and continue — via the AKAI's MIDI THRU Connection — to the DATUM, enabling the percussionist to affect the DATUM's program in some way.

Octapad Set-up

Each Octapad has eight rectangular surfaces, which can be struck with snare drum sticks to produce one of two types of MIDI message: Key On — Key Off, and Program Change, depending on whether or not a Program Change pedal is activated when a pad is struck.

Every time a pad sends a Key On message on MIDI channel 1 to the DATUM, the Note Number in that message becomes the new seed-value of SR which combines with FB to produce Polynomials, as explained above. Its MIDI Key Velocity, when it is not equal to zero, also becomes the new value of Key Velocity for all notes produced by the DATUM.

Every time a pad sends a Program Change to the DATUM, the Program Number selects either one of four values of FB causing a change in the ostinato length and its melodic contour, or one of four tempo presets causing the DATUM to play slower or faster. The values assigned to these Program Numbers were arbitrarily chosen in consultation with Graeme and are shown in Figure 7.

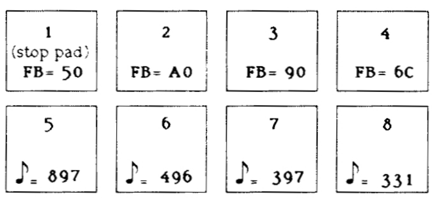

Figure 7

The tempo of the DATUM is set without the use of MIDI clocks. A delay subroutine, called many times for every note, sets the duration between transmission of the DATUM's Key On and Key Off messages. Values selected by the Program Change message determine the number of repetitions of this delay[7].

Because a single Octapad cannot send Program Change messages on two MIDI channels, it was necessary to use two Octapads, one to select voices on the AKAI independently of selecting Ostinato or Tempo changes on the DATUM. Figure 8 gives the source and destination of all MIDI messages in *Monophonic Variations*.

	AKAI S900 (in)		DATUM (in)	
	MIDI message	MIDI channel	MIDI message	MIDI channel
Octapad 1	keyOn – keyOff	All	keyOn – keyOff	1st
	Program Change	1st	Program Change	no effect
Octapad 2	keyOn – keyOff	All	keyOn – keyOff	1st
	Program Change	no effect	Program Change	2nd
DATUM (out)	keyOn – keyOff only	2nd	no effect	

Figure 8

The Note Number for each pad is preset and the choice is made partly with an ear to the sound made on the AKAI, partly with an ear to the melodic contour produced by the DATUM (which also plays on the AKAI)[8].

In addition to the Program assignments in Figure 7 the top left pad of the Octapad is recognised by the DATUM as a 'stop' pad. Hitting it cues the DATUM to cease playing notes, until further cued by the percussionist. This will happen automatically on the next key stroke recognised by the DATUM.

The percussionist, as well as providing input for other aspects of the DATUM's operation, is thereby able to stop and restart it; an operation which puts the computer entirely under the control of standard drumming technique. This, in the context of the system's responsiveness creates an element of live theatre, as if an invisible ensemble — a consort of gremlins — is taking part in an improvisation with the soloist.

AKAI Key Groups

Sounds made by the AKAI are arranged in key groups. Some of these key groups produce a single sound polyphonically, which is accessible over the entire range of the keyboard. Others are set up as split keyboards producing different sounds in different regions of the keyboard.

The key group played by the DATUM was organised with the prepared piano in mind. The keyboard splits were organised to include pockets of mute keys, as though wire-cutters were used as a part of the piano preparation. This allows the DATUM to produce a rest whenever it sends a MIDI Note Number that is a mute key.

Rhythms played by the DATUM are derived both from this procedure, and from the fact that some of the Note Numbers played — at any of the four selectable tempi — are too low to actuate sound from the Sampler.

A possible assignment of sounds and mute keys is shown in Figure 9 and the effect this keyboard map would have on the stream of MIDI Note Numbers is shown in Figure 10.

Voice Name	Lower Limit (dec.)		Upper Limit (dec.)	
Mbira	56	A♭3	64	E4
Glock Trem.	45	A♭2	60	C4
Rim Shot	42	F♯2	49	C♯3
Snare Drum	82	B♭5	88	E6
Xylo F5	99	E♭7	102	F♯7
Cowbell	50	D3	52	E3
Bass Drum	40	E2	43	A2
BC Bowl	68	A♭4	75	E♭5
M2	90	F♯6	94	B♭6
BD	100	E7	103	G7
Marimba	32	A♭1	35	B1
Crtale F♯	28	E1	30	F♯1
Tube	33	A1	35	B1
S3	97	C♯7	100	E7

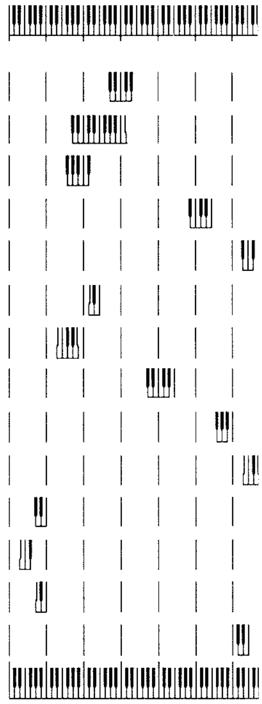

Figure 9

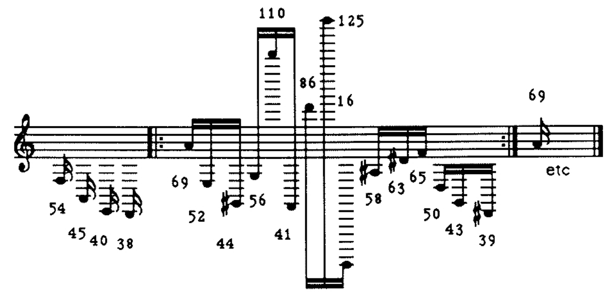

Melodic Contour produced by the DATUM is heard as Percussion ensemble played on AKAI Key Group 1

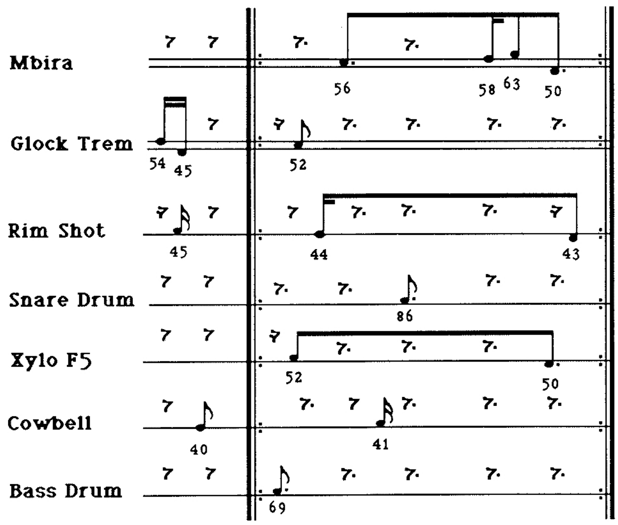

Figure 10

Controlling MIDI Feedback

The circular MIDI configuration in Figure 6 is essential for the multiple simultaneous conversations that take place in *Monophonic Variations*.

However, this configuration has the potential for MIDI produced by the DATUM to feedback, and this must be controlled to use it interactively[9].

The problem is best explained by an analogy — a global delay sometimes occurs while phoning ISD causing a speaker to be interrupted by his own voice. Likewise, the configuration in Figure 6 has the potential for the DATUM to interrupt itself with every message it sends.

Appendix 5 contains a further explanation of this process.

Aesthetic of Feedback

Interactive MIDI as it is found in *Monophonic Variations* involves communication between MIDI devices similar to other digital communications devices which are said to be operating in Full-Duplex mode. Full-Duplex operation involves a simultaneous two-way exchange of information. The closest analogy is the simultaneous monologue that takes place in plays by Albee and Beckett.

In MIDI this requires programming that drives MIDI as fast as it will go, attaining the full MIDI Bandwidth under all conditions, so that a performer is guaranteed an 'instantaneous' response from the instrument. The objective of such a response is a work that is performer-driven rather than system-driven[10].

If there is no such instantaneous response, the dialogue would only be Half-Duplex, which is analogous to conversational dialogue where one party listens while the other speaks.

By its nature, 'free' improvisation seems to have more in common with Full-Duplex communication — in both directions simultaneously — rather than Half-Duplex communication, which relates more to an exchange of words in conversation, where a delayed response is necessary for language processing. Improvised, or spontaneously composed music is a process involving instantaneous feedback and has more to do with exploring musical memory hidden within our biological reflexes.

Monophonic Variations creates a more human interaction with part of our environment by means of a self-regulatory process, called a Servo-system. A Servo-system operates on the basis of information it creates and feeds back to its input. It works on cues which it creates for itself, exhibiting a property of all self-regulatory systems which I call 'Servo-time'.

Servo-time is duration defined by the cues generated from within a process as opposed to Clock-time which is duration imposed externally onto a process without continuous reference to it. Servo-time is an attribute of a living system as opposed to a mechanical one, and accounts for the relationship between the computer and the performer in this work. This two-way feedback between performer and machine, as defined by the constraints of an interactive program in *Monophonic Variations*, has certain ramifications for me. Lines that define the musical caste imposed upon us by tradition begin to disappear. Roles change; the percussionist 'composes' on the spot; the composer becomes an 'absentee' performer; musicians manually take part in recreating the instrument to neutralise the artistic strategies hatched in some boardroom.

Footnotes

1. Gammatron are now manufacturing a 6809 version of the DATUM which will be known as the MIDI Tool Box.
2. I am indebted for Ian Shanahan's contribution to my understanding of Polynomials.
3. I am indebted to an article by Bernie Hutchins in **Electronotes** vol.8 #64 (April 1976), and Carl Vine, whose machine called the Patent Little Marvel, first introduced me to this method of sequence generation.
4. 7F Hex (01111111) acts like a mask where '1' corresponds to a hole that makes the bit behind it visible. A '0' in the mask makes the bit behind it invisible, resulting in a '0' for that bit position.
5. Exclusive OR outputs a '1' if both inputs are not the same.
6. The feed-back generated in the transient part of the sequence does not behave precisely according to the hardware model shown in Figure 3, which only applies to the periodic part of each sequence.
7. For example, Program Number 4 will repeat the delay 7 times. Each delay is ten milli-seconds long, ie. 1/100th of a second long. The note are just over 70 milliseconds long (7 or 7 x 10, plus 960 micro-seconds), or a tempo of approximately 14 notes per second.
8. In addition to these pads, extra pads were made from six miniature loudspeakers wired up via balanced sockets to Octapad 1. These, when struck, behave like a normal pad, giving a total of 22 pads available at one time 8 + 6 on Octapad 1 and 8 on Octapad 2. Each pad, Roland or 'cheap-skate', can have its own MIDI Note Number assigned, four times in four different patches.
9. The use of the term feedback here should not to be confused with the digital feedback process involved in the generation of Polynomials.
10. To achieve the maximum MIDI bandwidth, only the MIDI input is interrupt driven. MIDI output is polled and all processing associated with the piece is organised into sections that can be done 'on the fly', ie. between the moment a MIDI byte is first dispatched to the MIDI output register, to the moment it has finished sending it. This is the time taken for MIDI to send one start bit, eight data bits and one stop bit — 320 micro-seconds, ie. 320 x 1/1,000,000 of a second. The DATUM, even with a machine cycle time of 1 microsecond, can execute a sizeable block of code in that time.

Appendices

1.

BINARY	HEX	BINARY	HEX
0000	0	0001	1
0010	2	0011	3
0100	4	0101	5
0110	6	0111	7
1000	8	1001	9
1010	A	1011	B
1100	C	1101	D
1110	E	1111	F

2.

FB=A0 SR=24

Step	Carry	REGISTER	MIDI Note Num. Hex	MIDI Note Num. Dec
1	1	92	36	54
2	0	49	2D	45
3	1	A4	48	72
4	1	D2	76	118
5	0	69	2D	45
6	0	34	38	56
7	0	1A	3E	62
8	1	8D	31	49
9	0	46	2A	42
10	1	A3	47	71
11	1	D1	75	117
12	0	68	2C	44
etc	0	34	38	56

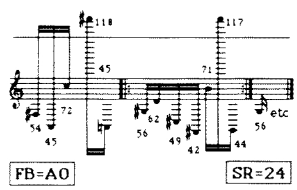

Figure 1

FB=A0 SR=38

Step	Carry	REGISTER	MIDI Note Num. Hex	Dec
1	1	9C	40	64
2	0	4E	32	50
3	1	A7	4B	75
4	1	D3	77	119
5	0	69	2D	45
6	0	34	38	56
7	0	1A	3E	62
8	1	8D	31	49
9	0	46	2A	42
10	1	A3	47	71
11	1	D1	75	117
12	0	68	2C	44
etc	0	34	38	56

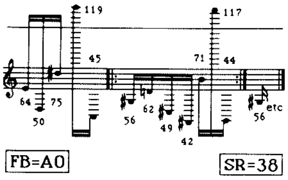

Figure 2

FB=A0 SR=44

Step	Carry	REGISTER	MIDI Note Num. Hex	Dec
1	0	22	26	38
2	0	11	35	53
3	1	88	2C	44
4	0	44	28	40
5	1	A2	46	70
6	1	D1	75	117
7	0	68	2C	44
8	0	34	38	56
9	0	1A	3E	62
10	1	8D	31	49
11	0	46	2A	42
12	1	A3	47	71
etc	1	D1	75	117

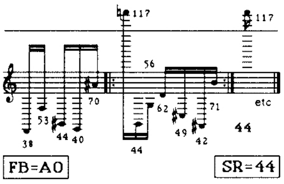

Figure 3

FB=50 SR=24

Step	Carry	REGISTER	MIDI Note Num. Hex	Dec
1	0	12	36	54
2	0	09	2D	45
3	1	84	28	40
4	1	C2	66	102
5	0	61	25	37
6	0	30	34	52
7	0	18	3C	60
8	0	0C	30	48
9	1	86	2A	42
10	1	C3	67	103
etc	0	61	25	37

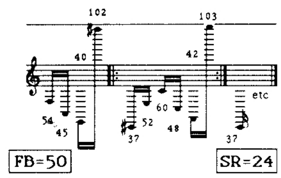

Figure 4

FB=50 SR=38

Step	Carry	REGISTER	MIDI Note Num. Hex	Dec
1	1	9C	40	64
2	0	4E	32	50
3	0	27	2B	43
4	1	93	37	55
5	0	49	2D	45
6	0	24	28	40
7	1	92	36	54
etc	0	49	2d	45

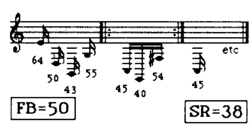

Figure 5

FB=50 SR=44

Step	Carry	REGISTER	MIDI Note Num. Hex	Dec
1	1	A2	46	70
2	1	D1	75	117
3	1	E8	0C	12
4	0	74	38	56
5	1	BA	5E	94
6	0	5D	41	65
7	1	AE	52	82
8	1	D7	7B	123
9	1	EB	0F	15
10	0	75	39	57
etc	1	BA	5E	94

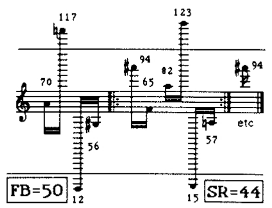

Figure 6

FB=90 SR=24

Step	Carry	REGISTER	MIDI Note Num. Hex	Dec
1	0	12	36	54
2	0	09	2D	45
3	1	84	28	40
4	0	42	26	38
5	1	A1	45	69
6	0	50	34	52
7	0	28	2C	44
8	1	94	38	56
9	1	CA	6E	110
10	1	65	29	41
11	1	B2	56	86
12	1	D9	7D	125
13	1	EC	10	16
14	0	76	3A	58
15	0	3B	3F	63
16	0	1D	41	65
17	0	0E	32	50
18	1	87	2B	43
19	0	43	27	39
etc	1	A1	45	69

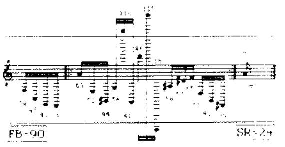

Figure 7

FB=90 SR=38

Step	Carry	REGISTER	MIDI Note Num. Hex	Dec
1	1	9C	40	64
2	1	CE	72	114
3	0	67	2B	43
4	1	B3	57	87
5	1	D9	7D	125
6	1	EC	10	16
7	0	76	3A	58
8	0	3B	3F	63
9	0	1D	41	65
10	0	0E	32	50
11	1	87	2B	43
12	0	43	27	39
13	1	A1	45	69
14	0	50	34	52
15	0	28	2C	44
16	1	94	38	56
17	1	CA	6E	110
18	0	65	29	41
19	1	B2	56	86
etc	1	D9	7D	125

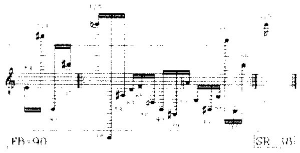

Figure 8

FB=90 SR=44

Step	Carry	REGISTER	MIDI Note Num. Hex	Dec
1	0	22	26	38
2	1	91	35	53
3	1	C8	6C	108
4	0	64	28	40
5	1	B2	56	86
6	1	D9	7D	125
7	1	EC	10	16
8	0	76	3A	58
9	0	3B	3F	63
10	0	1D	41	65
11	0	0E	32	50
12	1	87	2B	43
13	0	43	27	39
14	1	A1	45	69
15	0	50	34	52
16	0	28	2C	44
17	1	94	38	56
18	1	CA	6E	110
19	0	65	29	41
etc	1	B2	56	86

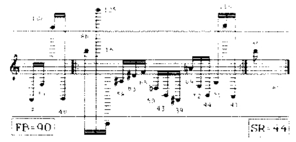

Figure 9

3. Here is a listing of the Polynomial Generator in 6802 Assembler Mnemonics, together with comments. Though this routine has since been tightened up considerably, it appears in the form that I used in *Monophonic Variations*.

Polynomial Generator

```
C6 08           LDA B #08       set bit tally
A6 01           LDA A $01,X     load SR
A4 02           AND A $02,X     AND it with FB
A7 03           STA A $03,X     copy to Parallel Input buffer

0D      shift   SEC             tie loose input to 1
A6 04           LDA A $04,X     load Serial Input buffer
49              ROL A           rotate left
A8 03           EOR A $03,X     Xor this bit with PI
A7 04           STA A $04,X     copy to SI
5A              DEC B           reduce bit tally
26 F5           BNE shift       end bit ? good! Feed-back bit ready

A6 01           LDA A $01,X     now get SR
46              ROR A           and shift it in
A7 01           STA A $01,X     copy new bit pattern
```

The first part of this routine gets a 'snap-shot' of the contents of the Shift Register (SR) through the 'holes' in the Feed-back Mask (FB). This is held as a value and applied to the Parallel Inputs of the XOR gates (PI), while the logical condition of each XOR gate is being generated. The second part determines the logical condition of the final XOR gate, by rotating each XOR gate left into the next XOR input eight times.
The third part of this routine shifts the Carry bit right, which now represents the Feed-back condition, into the Shift Register thereby generating a new eight bit pattern. This program generates a value in less time than it takes to transmit one MIDI byte, ie. 320 microseconds. All the processing required to generate a Polynomial and Filter it can commence as soon as the MIDI Status Byte is transmitted, and still have data ready before the MIDI hardware signals that it is ready to send the first MIDI Data byte.

4. Here is the program for filtering MIDI Note Numbers in *Monophonic Variations*.

MIDI Note Number Filter

A1 09	bandpass	CMP A $09, X	compare value with 20_{16}
2D 04		BLT transpose	it's less !
A0 09		SUB A $09, X	must be greater, reduce by 20_{16}
20 F8		BRA bandpass	try again
AB 0B	transpose	ADD A $0B, X	add 24_{16}
84 7F		AND A #7F	make sure it's a MIDI Data Byte
39		RTS	done !

When the eight-bit numbers produced by the Polynomial function generator are filtered, the MIDI Filter must distinguish between:
(a) those numbers from 2016 to 7F16 (all positive) and
(b) those below 20 Hex including those below 00 Hex (ie. all the negative numbers from FF Hex (-1) to 80 Hex (-128).
A modulo 20 Hex operation is performed on the numbers in the first category, which restricts the band to a range of 20 Hex (ie. 32 semitones).
Then 24 Hex — i.e. 36 — is added to each of the results in both categories, transposing the band up 36 semitones, and simply increasing the value of the other numbers.

Finally, all eight-bit results are converted to seven-bit results to comply with the MIDI specification for MIDI Note Numbers. This is accomplished with a mask value of 7F Hex.
5. Here is the listing of the MIDI ECHO sub-routine.

MIDI Echo Subroutine

37	PSH B	save the old B
F6 40 01	LDA B $4001	flush MIDI data receiver
0F	SEI	put the others to sleep
B7 40 01	STA A $4001	say hello in MIDI
F6 40 00 echo	LDA B $4000	Get MIDI receiver status
57	ASR B	try a bit
24 FA	BCC echo	is something there ?
33	PUL B	Good ! Put the old B back
0E	CLI	Wakey ! Wakey !
RTS	RTS	Bye

The DATUM turns off the interrupt system, and clears the Receiver Status before it sends a MIDI message. It then loops on MIDI receiver status, and once flagged by Receiver Ready, turns the interrupt system back on. This guarantees that the DATUM doesn't interrupt itself, and only genuine (i.e. performer driven) interrupts get serviced.

Notes on sociological and methodological stance with respect to music

Mark Rudolph

First, let me be clear that I believe the music which has the greatest potential of being essentially "new" music, i.e. music uniquely characteristic of this time and the immediate future, is music processed and synthesised by numerical algorithms. More specifically, it is music created by the process of deconstructing and coding chosen sound events into characteristic signal information, and the subsequent reconstruction of new events which inherit specific sound characteristics and musical expression from the choice and mix of materials from an accumulated collection of coded events. It is this use of computation in music which I am interested in.

It would, however, be better if there were no terms such as "computer" music (or even "new" music for that matter). The term music suffices to describe what we are talking about (and doing - music also speaks for itself). Music which is synthesised from analysed information is simply music with a somewhat different history of production than music produced by live instruments. In every case music is sound which we pay attention to, and hence give some significance to (at least this). The origin of the sound is mainly irrelevant to its existence, the origin being an untraceable set of influences and intentions receding into what we charmingly refer to as "the past". The use of digital analysis, processing and synthesis allows one to develop a musical expression which is fresh and intriguing, but the use of technology in music does not legitimise the music produced. The music legitimises itself (or does not) and obviates any explanation as to how it was produced, with what equipment, what algorithms, who produced it - ideology, what colour shirt was worn, etc.

However, despite the untraceable origin of music, it is important (and unavoidable) to adopt a stance with regard to the social circumstances in which we hear, perform or talk about it. Even the

category in which music recordings are placed for sale in a store can make us more or less receptive and attentive to it, or cause us to ignore it completely. The kind of music which we are capable of producing today (and are not producing consistently) cannot be categorised conveniently. This music should not be thought associated only with an institutional or intellectual environment, neither should it be associated solely with the commercial music industry. These two groups are virtually the only ones using digital equipment, probably due to the high cost. In almost all cases the equipment purchased consists of black boxes containing the same cheap mass-produced integrated circuits, configured in slightly different ways and sold for profit by one of a few large consumer electronics corporations. Frequently, the equipment is used by those who do not program their own software environment and tools, and who do not contribute to the construction of any specialised digital music hardware. What is clear is that to produce music of our time requires not only musical sensitivity, skill and independence, but also conceptual understanding gained by personal experience with algorithms and software. Work in a personally designed digital signal processing environment should result in a sort of "virtuosity" in the new techniques.

As well as discussing a sociological stance, I would like to outline what I regard as a particularly promising methodology for music generation, which can be referred to as "synthesis from reduced performance information". How the techniques associated with this method are used for musical expression is best left to musical examples such as the author's speech/music piece "Beautiful but marred by the blemish of a Perpetual Dissatisfaction" on the accompanying NMATAPE.

The piece consists of sound/music formed by synthesis from information acquired through the reconstruction of vocal poetic "performances". These textual "performances" were first analysed by the composer's program "K" which, using mainly enhanced methods of linear prediction, extracted time-varying sequences of independently modifiable acoustic information which characterises the musical and dramatic expressiveness of the particular vocalisation. Various combinations and modifications of information were specified and provided to the composer's program "V" which then synthesised the sounds abstracted from the analysed information.

The work contains passages which correspond to partial vocal information yielding environmental timbres or pure speech contours and energy contours. In addition, the energy flow, pitch contour and even the rate of passage of time are sometimes changed in order to produce an altered or enhanced dramatic effect. Finally, variations in the persona of the speaker such as female pitch range, old age or murmurings of a baby, create an implied "narrative".

What can be said in the context of this brief article is that first, if much of the meaning in music (as in speech) is in the inflections and intentions given by the precise movements of muscles and tendons, the quiver of the voice etc., then a method of representing sound events in terms of information must be specific enough to exactly characterise "performance" events. Second, the method must be universal, i.e. **any** sound must be adequately represented by sequences of simple standard structures of information parameters. Finally, the nature of the information must be such that the groups of parameters are de-correlated, i.e. that changes of one parameter do not mask the characterising features of other parameters.

Thus, what is needed is an informational representation of sound events which captures the essential expression and specific articulation of each event, yet is also a universal representation,

and has the property that its parameter groups are orthogonal, i.e. independently modifiable with combinatorial effect. If every sound event can be represented in a common information structure which acoustically characterises the event, then the source of the information is irrelevant when that information is used to re-construct new sound events. In addition, it is usually much easier to modify standard and orthogonal information parameters to achieve particular expressive effects than it is to modify recorded sound directly. New events may be composed by combinations of information streams, modulation of one stream by another, or by abstraction of streams from mathematical functions. Many musically significant possibilities arise when the modifications of parameters are correlated to or combined with other information sequences stored in a library of sound performance characterisations.

Note that the level of information which is critical to this methodology is more "fine grained" than the note level information associated with MIDI. The goal of the method is to represent sound events by sequences of standard form structures of decorrelated, and hence independently modifiable, information parameters capable of reproducing the unique event itself. In contrast, MIDI represents an entire class of equivalent sound events which must be realised on some pre-configured hardware. A fundamental consequence of the representation of sound events by sequences of independently modifiable acoustically characteristic information parameters is that new sound events may be abstractly constructed combinatorially or time-varyingly according to desired acoustic features and specific articulation and expressive content. The synthesis from reduced performance information can be performed in real-time using multiprocessor DSP hardware and new physical articulation controllers (WX7 windcontroller, multidimensional keyboards, etc), or even conventional instruments fitted with special switches and analysis hardware. The method is intended to facilitate the creation of expressive music by means which have never been available before. The "playing" of the characteristic information which is transformationally equivalent to unique sound events may be either origin or consequence.

Computers — performance of, and with

Compositional techniques

Cindy John

In this article I will discuss the compositional procedure for my works *Blowout* and *Death of an Insect*. Work in progress at present involving computerised vocal sounds will also be discussed.

Blowout is a composition written for six trombones and computer. In this work the computer part was written first, using the trumpet algorithm in the music 4BF program on the Universe68 computer. Apart from the usual parameters entered for each note, such as start time, duration , amplitude and pitch, this particular algorithm offers at least sixteen extra parameters providing unlimited scope for finding different timbral effects. Because the particular computer in use does not permit real time operation, I made approximately 30 short scores testing the effects of each parameter when altered from the normal array emulating a trumpet sound. A comparison of the parameters is shown. Extensive changes in the sound of this score were especially affected by the parameters marked. After the selected timbre had been chosen, work on a pitch structure for the music, which would also be used for the trombone parts, began. The pitch frame is based on semitones and tritones, being compatible with the harmonic boundaries of the trombone. With these ideas in mind, two contrasting types of computer sound were devised. The first consisted of compressed, sharp timbres gradually increasing in tempo and length, and the second of long slow evolving textures bringing out colours in other ways. Once this was completed, and a multi-track version finished, the timings of the computer part were notated on manuscript in graphic form. At this point the 6 trombone parts were written against the computer score. Minor 2nds, major 7ths and tritones form the opening trombone passage. As the tempo of the computer part increases, the accompanying trombone parts are overlaid with longer durations, with particular phrases being brought out of the texture. This sec-

Normal		Altered for timbral change
4		4
0		0
3		3
0.7		7
9.07	fractional pitch deviation	9.0762
0		0
7		7
0.02		0.02
0.15		0.19
0.033		0.09
0.06		0.09
0.06	Centre frequency	0.09
8	deviation function no.	11
8.04		8.04
1		1.7
3.523		4.942
0	peak deviation of vibrato in	0
0.33	% of cent.frq.	7
7	rate of vibrato	30
0.02		0.09
0.01		0.0
0.5		0.5

tion culminates in a texture of clustered semitones in semiquaver form, finally opening out with expanded flourishes from the trombones.

The second section opens with a higher-pitched entry from the computer, using a slightly different timbre but similar structure to the first part. Here the trombones reflect the computer part with chordal entries gradually changing into fast executions of quintuplets. Again with overlays, it builds to a climax, this time ending with the computer part.

The third section begins with the trombone gesture arriving before the computer. The trombones form a very thick chordal texture which competes with the computer, ending at the same time for the final climax. After this section the mood changes. There is no strict pitch frame for the trombones here. Their parts were written whilst listening to the computer part and with the harmon mute in mind. Each computer timbre reflected the melodic and harmonic phrases of the trombones. The mute was used differently for each part, making a continuous flowing pattern of intended sombre sounds.

The work attempts to produce a tension between the trombone and computer parts, climaxing in the centre of the piece and showing reflections of the previous mood in the later half. The parts were combined in the recording studio[1]. Each trombone part was recorded on a multitrack recorder in sections which were later mixed with the computer part. Additional effects, including reverberation and equalisation, were added in the final stages of the mix. The title *Blowout* refers to the idea of tension and release, as well as the vision of energetic trombone players. This piece was written and revised during 1987-88.

Death of an Insect is a work based on a poem by Charles Reznikoff. The following text has been set for computer and voice:

The sparrow with its beak taps the beetle
and it begins to buzz loudly
as if the bird has set off an alarm clock.

The beetle flies into the air
in a series of clumsy gyrations
and the sparrow follows it gracefully.

In this piece the computer part was generated using two instrumental algorithms: piano and strings. The piano algorithm was used without modification to its only extra parameter, the function envelope. On the other hand, the string algorithm, which offers many alternatives to the sound texture, was used both modified and unmodified. The piece commences with a manipulated version of the strings, giving a sense of flight and power relating to the sparrow in the poem. In the following section a recorded version of my own voice introduces the text. The voice was not computerised but treated with a pitch shifter and digitised reverberation in the studio. A music concrete sound, comprising a large part of the third section, was obtained from an aluminium bowl filled with water. Speed manipulations and overlayering on the multitrack provided a gradual heightening in the texture of the sound. Small elements of the final section were created using the concrete material. The aim was to depict an atmosphere of the space in which insects fly, and to describe the erratic movements of injured insects. The final section of the piece portrays the hopelessness of recovery once death is finally at hand. The extra parameters were not used in this section, only the string and piano algorithms. The pitch structure of this piece revolves around a freely composed vocal line. Therefore, the string and piano parts at the end of the piece are extensions of the vocal line, at times in canonic form.

Death of an Insect has two versions; one as described above, and another version which uses live voice and the computer in the final section alone. The instrumentation for version 2 is:

Computer — on tape
High Voice- with amplification
Piano
Percussion: Water Gong — with amplification
Sand Blocks — with amplification
Wood Blocks
Ratchet
Timpani

The structure of this version is slightly different, making the piece shorter than the other, 5 minute version. Both were written in 1988.

Presently, I am directing my composition towards performers and the computer, including the computerised voice. It is now possible to write computerised vocal sounds which are processed on the Universe68 computer; Jim Sosnin and Graeme Gerrard got the program on the road again. So far, I have only made preliminary test scores to see what composing tools are available. The effects of overlayering and looping of segments of sampled sounds are of special interest. Added to this facility is the use of the recorded voice in the studio and naturally, the live voice itself. Compositions underway at present are *Extentions7*, *5 Songs*, and an as yet untitled work for chamber ensemble and computer.

Extensions7 is made up of 7 short pieces for prepared piano based on a pentad. Each piece is being extended using the computer algorithms and the voice sampling program. The 5 Songs are also being written for computerised voice even though they stand alone as two part songs.

The untitled chamber work has most of the computer part written, but this will also be extended

with sampled voice. This work is based on 2 matrix patterns. Each matrix is divided up into symmetrical parts from which pitch groups are selected. These groups are used in linear and vertical methods, each group following a chosen rhythmic frame. At some stages the matrices cross paths making way for larger selective pitch frames. The chamber instruments as well as the sampled voice will use the above method for the score.

In these compositions it has sometimes been appropriate to write the computer part first and at other times, the opposite. This decision has been governed by the textural result when they are drawn together. Pitch frames and rhythmic gestures, after initial creation, are sometimes changed as the works progress. Although this way of writing with technological facilities may be considered slow, I find it most suitable for my compositional ideas and working patterns.

Footnote

1. My thanks to Simone de Haan for the performance and Jim Sosnin of the La Trobe University music department for technical assistance.

A digital signal processor for the real-time processing of sound

C. H. Dick

Introduction

Digital signal processing offers the audio researcher many powerful techniques for sound synthesis, analysis, and processing. Although the processing may be done in an off-line fashion, there are many instances where real-time operation is desirable. A real-time DSP system could implement instruments that can be used in a performance setting, and that are limited in their sonic capabilities, and **playing** technique, only by the programmer's imagination. A general purpose system means that virtually any type of sound processing technique can be implemented with the one piece of hardware. Future audio processing techniques could be implemented without the need to develop more hardware, only the system software need be changed.

The enormous potential of DSP techniques for audio processing, coupled with the recent availability of low-cost DSP microprocessors, has paved the way for the development of economical digital audio processing systems. Although DSP microprocessors offer considerable arithmetic capacity, there are applications where the processing power of a single DSP chip is insufficient. The approach taken in this project was to implement an array of DSP microprocessors to produce a high performance audio processing system.

This paper describes the multiprocessor hardware architecture — providing details of the system DSP process modules, peripheral devices and host interface. An overview of the system software includes a description of a TMS32020 monitor program, multiprocessor operating system, and application software development tools. Finally, an application example is presented which shows how the system may be programmed to implement real-time cross-synthesis.

Hardware Architecture

The basic building block of the multiprocessor system is a general-purpose signal processing card based on the Texas Instruments TMS32020 microprocessor. The card supports various types of memory, provides communication to other processing cards via a high-speed synchronous parallel bus and a 2.5 Mbit/sec synchronous serial link, and supports interfacing to peripheral devices such as ADC's, DAC's and asynchronous serial communication circuits via a 16-bit parallel IO bus. A simplified version of this card, referred to as a processing module (PM), is shown in figure 1. A card loaded with a TMS32020 is capable of 5 MIPS., with the TMS320C25 this is increased to 10 MIPS.

The signal processing card developed for this project provides 20K of zero wait-state RAM that can be allocated in 4K blocks to either the processors data space or program space by inserting/removing option jumpers. A 4K block of ROM, configurable as zero, one, or two wait-state memory is mapped into the bottom of the program address space. A memory expansion bus is supported by the DSP card. The buffered address,

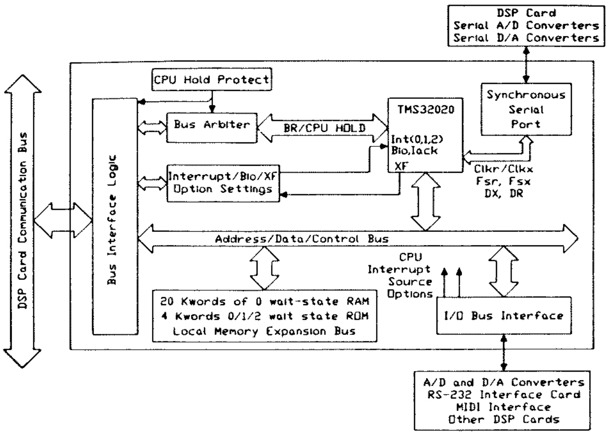

Figure 1. TMS32020 based DSP Processing Module (PM).

data and control signals are available at a 96-pin eurocard style connector, and additional data or program memory can be added to the system by plugging a memory daughter board onto the memory expansion bus connector.

In addition to the parallel bus over which processor cards can communicate, transferring frames of sampled data, parameter sets or any other data, a separate parallel bus is supported on the card for interfacing to peripheral devices.

The I/O bus takes the physical form of a 50-way header on the card that brings out the 16 bit data bus, port address bus, and control bus. The TMS32020 micro-state-complete output, has also been provided on the IO bus to allow the easy insertion of I/O wait states. A 50-way ribbon cable, with interleaved earth's, can be run from the I/O header on the processor card to a shielded enclosure containing the system I/O devices. In this manner noise sensitive devices such as DAC's and ADC's can be isolated from the noisy environment of the multiprocessor enclosure.

A single card can be partially loaded and run stand-alone for applications where the computing power of a single TMS32020 is sufficient. The processor card can also be used as the basic building block for larger signal processing systems. Based on this concept, a system capable of supporting 8 processing cards was implemented. This was accomplished by connecting the processing cards together over a high speed parallel backplane. At the time of writing 4 PM's have been constructed. The hardware associated with interprocessor communication was designed to support block memory transfers between processing cards so as to exploit the TMS32020's block move from data memory to data memory instruction, BLKD, and block move

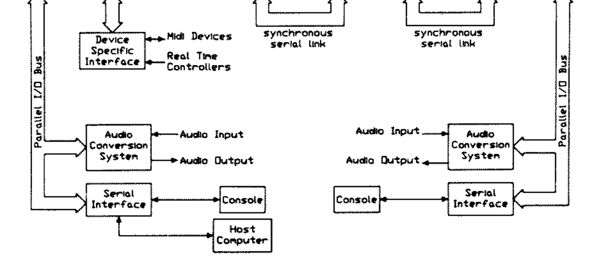

Figure 2. Multiprocessor System Hardware Overview.

from program memory to program memory instruction, BLKP. If the TMS32020 instruction repeat feature is coupled with the BLKD instruction for transferring data memory from one processing card to another, peak transfer rates of 5 Mwords/sec (80 Mbits/sec) can be achieved. This figure does not include overheads associated with obtaining use of the communication bus. Figure 2 gives an overview of the system.

Interprocessor Communication in an Array of DSP Cards.

Data is transferred from a source processing card to one or more target processing cards by halting the target PM's and gaining direct access to the off-chip local memory of the halted card(s). Transferring data to a PM generates an interrupt to the destination PM, so signalling the presence of new data or a system message to it. Buffer contention is handled using semaphores. The transfer operation may also be made invisible to the destination PM.

The system offers a **broadcast** mode of data transfer, whereby a source processing module may transfer data to any number of target modules in a single global bus operation.

External logic arbitrates for control of the global bus, asserting a bus grant signal when global access is allowed by the requesting processor. By latching the processor bus request line, BR, a processor can attain control of the global bus and perform block memory transfers. Latching the bus request signal is necessary in implementing a block memory transfer scheme, because the processor bus request output is removed while the processor is obtaining local data to be transferred to another processing module. If the bus request signal were not latched, several requests asserted simultaneously could cause wait state cycling of a processor in the middle of a memory block transfer. Servicing of bus requests is done on a prioritised basis. Each of the processing cards in the system are essentially identical, however each may be assigned a bus access priority by insertion of a jumper link. The core bus arbitration logic resides on a processing card, only one card in the system has this group of components loaded. The latched bus request signal of each processing module is buffered onto the system backplane, this group of 8 request signals form the input to the bus arbitration logic. The core element of the bus arbiter is a programmable logic array. The bus arbiter asserts one of its 8 bus grant lines in response to the pending requests. The 8 bus grant outputs of the arbitration logic are buffered onto the backplane. Each of the processing cards tap the bus grant signal associated with the bus request line that it is using off the backplane. This is then used locally on the card to generate a memory ready signal. While a processor is waiting for access to the global bus after submitting a bus request, wait states are inserted.

The bus arbitration logic can be considered as operating as a transfer-request queue. The bus request output of each card is clocked into the bus request latch on the rising edge of the processor CLKOUT1 clock. Logic connected to the bus-request-latch output detects the presence of any pending requests, and disables sampling of subsequent bus requests until all current requests have

been serviced. Once the current bus master has completed its data transfer across the global bus and removed its bus access request, the request with the next highest access priority is serviced. After all current requests are serviced, bus-request sampling is re-enabled. In this fashion, processing cards with a low bus access priority are guaranteed an uninterrupted time slice in which to perform data transfers to other PM's.

This scheme has its advantages and disadvantages. Having to explicitly specify a processor hold-code and clear a bus request incurs an instruction overhead for data transfers across the global bus. However, such an implementation allows for a broadcast mode of data transfer and permits uninterrupted transfers of blocks of data.

The mechanism for identifying the destination module in a data transfer is a based around a an 8 bit wide hold-identification data latch (hold-id latch). Once a PM has attained bus mastership, it writes a PM hold-code into the hold-id latch. One bit of the hold-id latch is assigned to the HOLD input of each TMS32020 resident in the system. Driving HOLD active causes the TMS32020 to halt instruction execution after the current instruction cycle is complete. Table 1. provides the timing details for a data transfer across the system bus.

Program Task	Time (Number of Cycles)
Attain Global Bus/ Halt Target PM's	12
Data Transfer	N X 7
Release Halted PM's/ Remove Bus Request	9

Notes: N = the number of 16 bit words transferred

Table 1. Global Bus Data-Transfer Timing.

A PM can hardware protect itself from being halted by another PM. The bus master can determine which PM's are halt-protected.

Each PM has the TMS32020 XF output flag and interrupt acknowledge signal buffered onto the system backplane. These may be jumpered to any of the interrupt/BIO inputs of other PM's.

Peripheral Devices

Analogue to Digital and Digital to Analogue Conversion

Each PM may have 2 analogue conversion systems connected its parallel I/O bus. The conversion system is implemented with 3 cards:
1) 12-bit ADC/16-bit DAC card,
2) input anti-aliasing card, and
3) output reconstruction filter card.

The ADC/DAC module contains the PM I/O bus interface logic as well as logic to buffer the ADC and DAC.

At the industry standard sampling frequency of 44.1 kHz, the TMS32020, with an instruction cycle time of 200 ns, can at best perform 113 instructions per sample period. This figure is reduced if multi-cycle instructions are used. Thus, to provide the programmer with some flexibility in trading signal bandwidth for processing complexity, a 16 bit timer, programmed by the controlling PM, is included on the conversion module and is used for generating the sample clock. Changing the sampling frequency also requires changing the input and output filter cards. At present 2 pairs of filter cards have been constructed: one with a cutoff frequency of 4.4 kHz, the other with a cutoff frequency of 8 kHz. All filter cards are 9th order elliptic designs with a rolloff of about 777 dB/decade.

The steep filter rolloff results in a very non-linear phase response. However, these filters were constructed for experiments using linear predictive coding techniques, and their choice was made out of consideration of information presented in references [3] and [6], below.

Work is currently in progress to implement a new conversion system based around a new Phillips device, the SAA7320, that offers a 256 times over-sampling filter and stereo 16 bit DAC in one package.

Asynchronous Serial Communication Interface

Up to two dual port RS232 asynchronous serial communication cards may be connected to a PM I/O bus. Each port has an individually settable baud rate. Connecting one port to a console was found useful during software debugging.

Application-Software Development Environment

The host computer to the multiprocessor machine is an IBM PC/AT. This choice of host processor was a consequence of the large number of TMS32020 software development tools available for the PC and its low cost. Program development tools available on the host include:
1) Texas Instruments cross-assembler,
2) a TMS320 software simulator,
3) Texas Instruments software development system (SWDS),
4) and a software interface to the multiprocessor system.

The interface software consists of two major components:
1) a monitor program run on the multiprocessor system, and
2) a small multiprocessor operating system running on the host.

TMS32020 Monitor Program

Each PM in the system runs its own copy of the monitor program from local ROM. The monitor services two RS232 serial data streams:
1) a link to a console, and
2) a host link.

Commands may be issued from the console to examine program and data memory, fill blocks of memory, move blocks of memory, exercise the PM communication bus and test the signal conversion modules. PM1 is the only PM connected to the host. Commands are issued from the host to PM1. Host commands not for PM1 are passed by PM1 to the appropriate PM using the parallel communication bus. The monitor software includes a library of routines that may be called from a user program. The library contains device drivers for the console and host serial link, as well as a range of I/O routines that are useful for program debugging.

SIG — an Operating System for the Multiprocessor

SIG is a program that provides the system user with a flexible interface to the DSP engine. In addition to implementing commands to initialise PM's, SIG provides an interactive mode of operation from which algorithm parameters may be **tuned** whilst the algorithm is running. Table 2 gives a brief summary of SIG's commands and illustrates the command syntax.

SIG Command	Description
pload [-n pmid] file	PM program load
dload [-n pmid] file -o blkstr	PM data memory load
pdump [-n pmid] -b blkstr -e blkend	PM program memory upload
ddump [-n pmid] -b blkstr -e blkend	PM data memory upload
peek [-n pmid] address	data memory examine
poke [-n pmid] address value	data memory deposit
system [pmid]	get parameter list
source file	redirect SIG command input
graphix	enter interactive mode
run pmid	run program on PM *pmid*
exit pmid	exit program on PM *pmid*
quit	return to host operating sys.
set par-name value	set value of parameter

pmid : process module identification number
file : file name
blkstr : start address of memory block
blkend : end address of memory block
value : hexadecimal number
par-name : application program parameter name
address : memory address

Table 2. SIG Command Language Summary.

System Software Initialisation

Before running an application on the DSP engine, each PM must have its program space, and possibly its data space initialised by the host. The details of initialising the PM array to perform a particular application are stored in the form of a SIG command file. This file contains all the program load and data load commands for each PM. The system is initialised by passing the appropriate command file name to the SIG **source** command.

SIG — Interactive Mode

SIG enables the user to interact with application programs by using the **system** command. In the TMS32020 application program source code, the names of program parameters that are to be made accessible from SIG are passed to a parameter table building macro. This causes the data structure shown in figure 3 to be built in the PM program space. An application program signals the presence of a parameter table by placing a table-present-code at a reserved location in program memory. If the table-present-code is found by the SIG **system** command, the parameter table is uploaded to the host. The parameter names are displayed in a spread-sheet format on the host console. The cursor keys are used to move between the value fields of the parameters, and simply typing a new value in the parameter value field causes a parameter change message to be sent to the appropriate PM. Each PM parameter table is allocated a full screen on the host console. Parameter screens are toggled by using console function keys.

parameter table present flag
parameter 1 name length field
parameter 1 name (max. 6 chars)
parameter 1 address
• • •
parameter n name length field
parameter n name
parameter n address
end parameter table marker

Figure 3. PM Parameter Table Format.

Writing an application program

The programmer must partition the desired algorithm into tasks for each PM in the system. Figure 4 shows a typical structure for a PM task.

system initialisation
interrupt 0: real time task ADC/DAC servicing
interrupt 1: used for message passing between PM's
interrupt 2: optional console link (PM1 uses this interrupt for host servicing)

Figure 4. Typical Organisation for a PM task.

LPC based Cross-Synthesis Implementation

Introduction

Figure 5 illustrates cross-synthesis implemented with linear predictive techniques. The basic idea is to filter a sound, x(n), with the resonance model of another sound, y(n). The signal x(n) is the excitation input to the time-varying digital synthesis filter, 1/A1(z). The resulting sound is intuitively predictable because of the perceptual relevance of excitation and resonance concepts. For example, if x(n) is a sampled piano sound, andj y(n) is a speech waveform, a type of talking piano effect is produced. Cross-synthesis can be implemented using other techniques, but the linear prediction realisation opens up many musical possibilities.

Linear prediction has been discussed extensively in the literature, ([1], [2], and [6]) and so a detailed treatment will not be presented here. Only

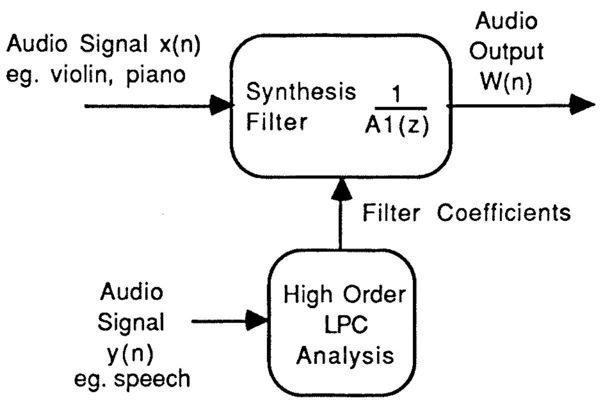

Figure 5. Cross-Synthesis implemented using Linear Prediction.

enough details to provide a suitable context for the results presented later, and for the following general discussion will be stated.

By viewing sound as the output of a resonant cavity driven by some excitation function, an LPC analysis applied to the system output effectively separates the excitation function parameters from the resonance cavity parameters. The parameters describing the excitation function are the pitch-period and gain. The resonant chamber is described by a set of filter coefficients. The analysis is applied to the signal at the parameter frame rate. The analysis thus produces a compressed, parametric representation of the sound. The usefulness of this representation is that the parameters map to physical aspects of the sound. The pitch period parameter relates to the original sound's pitch envelope, the gain parameter reflects the amplitude envelope of the sound, and the filter coefficients represent the resonance structure of the sound. The pitch and gain parameters can be modulated to achieve any desired re-shaping of the original sound's pitch and amplitude structure. In addition, the frequency and time aspects of a sound may be changed independently of one another: this is not always possible with other synthesis techniques, eg. sampling. The filter coefficients may be manipulated to shift the position of resonance peaks. However, this is sometimes difficult in practice as discussed by Moorer in [8]. Hybrid instruments can be made by interpolating between parameter sets of different instruments. The following sections discuss the details of implementing cross-synthesis on the multiprocessor.

Task Partitioning

The cross-synthesis algorithm has two distinct sections:

1) the LPC analysis, and
2) the synthesis.

Mapping the algorithm onto the machine architecture is straightforward. One PM is used to implement the analysis, and a second PM implements the synthesis. The analysis parameters are passed to the synthesis PM across the parallel bus.

Real-Time Fixed-point arithmetic LPC Analysis Implementation

The LPC analysis algorithm implemented on the TMS32020 is presented in figure 6. Calculations are performed using fractional fixed-point rounding arithmetic. Signal sampling is run under interrupts.

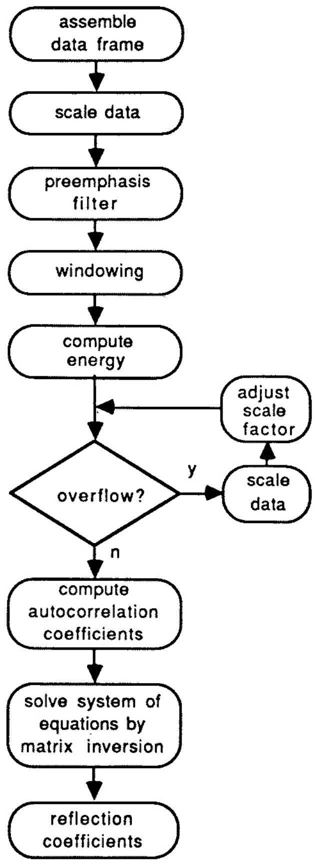

Figure 6. LPC Analysis Flowchart.

The data sequence is first scaled by a scale factor equal to the reciprocal of the current data frame amplitude peak. If this is not done the following windowing operation results in lost precision, since only the high order 16 bits of the windowed data sequence are kept. The data is then preemphasised and windowed. Application of the preemphasis filter decreases the system sensitivity to numerical errors [4]. The next step is to calculate the p+1 autocorrelation coefficients, where p is the order of the predictor. The autocorrelation coefficient at lag zero, r0, is the energy signal energy, and is guaranteed to be the largest autocorrelation coefficient. If calculation of r0 causes an overflow, the data sequence is scaled, and r0 computed again. This procedure is iterated until r0 can be calculated without overflow. The rest of the autocorrelation coefficients are then calculated. Although the TMS32020 operates on 16 bit data, the autocorrelation computation can be implemented in such a way as to retain full 32 bit precision for each autocorrelation calculation. This is done by having two copies of the data sequence, one in internal data memory, the other in internal program memory, and using the TMS32020 **MAC** instruction coupled with the instruction repeat feature. The 32 bit result is rounded and the high-order result kept.

Solution of Autocorrelation Equations

The autocorrelation equations are solved for the reflection coefficients:

$$k_m \qquad m = 1,2,\ldots,p$$

As shown in [5], the resulting synthesis filter will be stable if:

$$|k_m| < 1 \qquad m = 1,2,\ldots,p$$

This bound on the reflection coefficient magnitude is obviously desirable for fixed point implementations.

The autocorrelation matrix is symmetric and Toeplitz. Taking into account this special form of matrix, Levinson, Durbin, and Robinson have formulated efficient algorithms for solving this system of equations (see [9] for references to this work). These algorithms produce the desired reflection coefficients. However, as pointed out in [4] and [9], little is known about the range of the intermediate variables used in calculating the reflection coefficients. This causes bothersome scaling problems. The method finally implemented for solving the autocorrelation matrix is that formulated by Le Roux and Gueguen in [9]. This method has the desirable property that the magnitude of the intermediate variables used for calculating the reflection coefficients are bounded by unity, thus scaling is not a big problem.

Analysis Execution Time

Figure 7 is a plot of the LPC analysis execution time vs. predictor order, with analysis frame length as a parameter. The smaller analysis frame allows a larger order predictor to be implemented for a given sampling frequency, and minimises averaging of the time-varying signal at the expense of data compression ratio. This data was compiled by running the analysis on the TMS320C25 software simulator, for predictor orders of 4 to 52, in predictor order increments of 4. The times were calculated using an instruction cycle time of 200 ns, even though the TMS320C25 may be run with an instruction cycle time of 100 ns. The reason for this was to obtain time estimates for the TMS32020 processor. The exponential increase in execution time with predictor order is clearly discernible from the the graph. It is anticipated that increasing the sampling frequency would add a constant offset to the curves as the sample acquisition interrupt routine would be executed more frequently. No figures for this variation have been measured as yet. Table 3 provides a breakdown of the total analysis execution time for a predictor order of 28, 200 sample frame size and sampling frequency of 10 kHz.

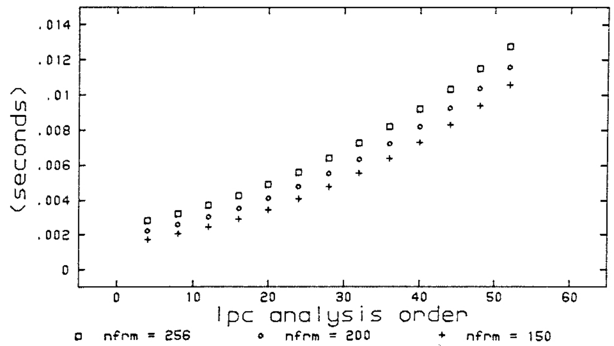

Figure 7. LPC Analysis Execution Time vs. Analysis Order. nfrm = number of samples in the analysis frame.

Fixed-point arithmetic Synthesis from LPC Parameters

The synthesis filter is implemented with a two multiplier lattice structure. Although this structure is expensive to implement in terms of the number of operations required (two multiplies, two additions and a data move), it offers several important advantages:

1) the reflection coefficients produced in the analysis can be used directly,
2) there is a simple test for filter stability,
3) the nature of the filter is such that it is less sensitive to numerical inaccuracies than other filter types, and
4) it is relatively easy to interpolate between filter coefficient sets.

A lattice filter section takes 2.2 microseconds of processing time on the TMS32020. Lattice filter sections are cascaded to implement a high order synthesis filter.

LPC Analysis Task	Time (μs)	% Real-Time
initial scaling	815	8.2
preemphasis filter	402	4.0
windowing	361	3.6
autocorrelation coefficients	1215	12.2
matrix solution	2192	21.9
data frame assembly	127	1.3
interrupt servicing	406	4.1
Total Time	5518	55.3

Notes: 1. sampling rate = 10kHz
2. parameter frame rate = 100 Hz
3. frame size = 200 samples
4. 50% frame overlap, new parameter frame every 10 ms
5. add 460 μs for each iteration through pre-autocorrelation calculation scaling loop
6. times are for TMS320C25, 200 ns instruction cycle time used in calculations

Table 3. LPC Analysis Timing Breakdown.

Discussion

In comparing the magnitude of the reflection coefficients computed with the TMS32020 fixed-point implementation, against those produced from a floating-point arithmetic implementation, differences are usually restricted to the 3rd decimal place.

In listening tests to date, no significant difference has been noticed between using the reflection coefficients produced by the fixed-point arithmetic analysis, and those produced from a floating-point arithmetic analysis, in the synthesis filter.

Several problems have been noticed with the cross-synthesis algorithm of figure 5. When y(n) is a speech signal, and x(n) is a musical instrument, the intelligibility of the words in the resultant synthesised signal is poor. This has been attributed to two main reasons:

1) the synthesis filter driving function is not spectrally flat, unlike the situation in standard LPC vocoders, and
2) the amplitude envelope of the synthesised waveform does not follow the amplitude envelope of the speech waveform.

Two approaches are being explored to overcome these problems. The spectral flattening nature of the LPC inverse synthesis filter is being used to whitten the signal x(n), and the resulting residual is used as the input to the synthesis filter. This is illustrated in figure 8. This cross-synthesis

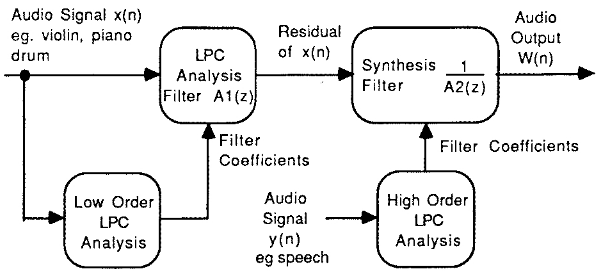

Figure 8. An improved Cross-synthesis system.

scheme is currently being implemented with 3 PM's: PM1 performs the low order LPC analysis and synthesis to produce the required residual, PM2 performs the high order LPC analysis, and PM3 implements the final synthesis using data from PM1 and PM2 that is communicated across the system bus.

A post-synthesis normalisation scheme is being explored to match the amplitude envelope of the speech signal to that of the synthesised waveform. The normalisation is being performed on a frame by frame basis.

Conclusion

The significant arithmetic capacity of the TMS32020/TMS320C25 make them suitable for implementing an audio processing system. The main drawback of this processor is the small word length, which is restrictive for some applications. However, useful audio processing can still be achieved with this processor. The newer TMS320C30, at 33 MFLOPS., would appear to be a step in the right direction for a future audio processing system.

From the results of above, the extended cross-synthesis algorithm can be implemented on the multiprocessor using 3 PM's.

Although the system was designed primarily for audio processing, it is general purpose in nature and can be used for processing of non-audio signals.

This paper was first presented at the Audio Engineering Society third regional convention, held in Melbourne in August 1988.

References

1. J. Makhoul, *'Linear prediction: a tutorial review'* **Proc. IEEE**, vol. 63, pp.561-580, Apr. 1975.
2. J. Makhoul, *'Stable and efficient lattice methods for linear prediction'* **IEEE Trans. Acoust., Speech, and Signal Processing**, vol. ASSP-25, pp.423-428, Oct. 1977.
3. J.D. Markel and A.H. Gray, Jr., *'A linear prediction vocoder simulation based upon the autocorrelation method'* **IEEE Trans. Acoust., Speech, and Signal Processing**, vol. ASSP-22 pp.124-134, Apr. 1974.
4. J.D. Markel and A.H. Gray Jr., *'Fixed-point truncation arithmetic implementation of a linear prediction autocorrelation vocoder'* **IEEE Trans. Acoust., Speech, and Signal Processing**, vol. ASSP-22, pp.273-282, Aug. 1974.
5. J.D. Markel and A. G. Gray, Jr., *'On autocorrelation equations as applied to speech analysis'* **IEEE Trans. Audio Electroacoust.**, vol. AU-21, pp.69-79, Apr. 1973.
6. J.D. Markel and A.H. Gray Jr., **Linear prediction of speech** (Springer, Berlin-Heidelberg, 1976).
7. R.E. Morley, A.M. Engebretson, J.G. Trotta, *'A multiprocessor digital signal processing system for real-time audio applications'* **IEEE Trans. Acoust., Speech, and Signal Processing**, vol. ASSP-34, pp.225-231, Apr. 1986.
8. J.A. Moorer, *'The use of linear prediction of speech in computer music applications'* **Journal of the Audio Engineering Society**, vol. 27, pp.134-140, Mar. 1979.
9. J. Le Roux and C. Gueguen, *'A fixed point computation of the partial correlation coefficients'* **IEEE Trans. Acoust., Speech and Signal Processing** June 1977, pp.257-259.
10. J.M. Snell, *'Professional real-time signal processor for synthesis, sampling, mixing, and recording'* **Audio Engineering Society preprint**, presented at the 83rd AES convention, New York, Oct. 1987.

Figure 9. Circuit board layout for a DSP module.

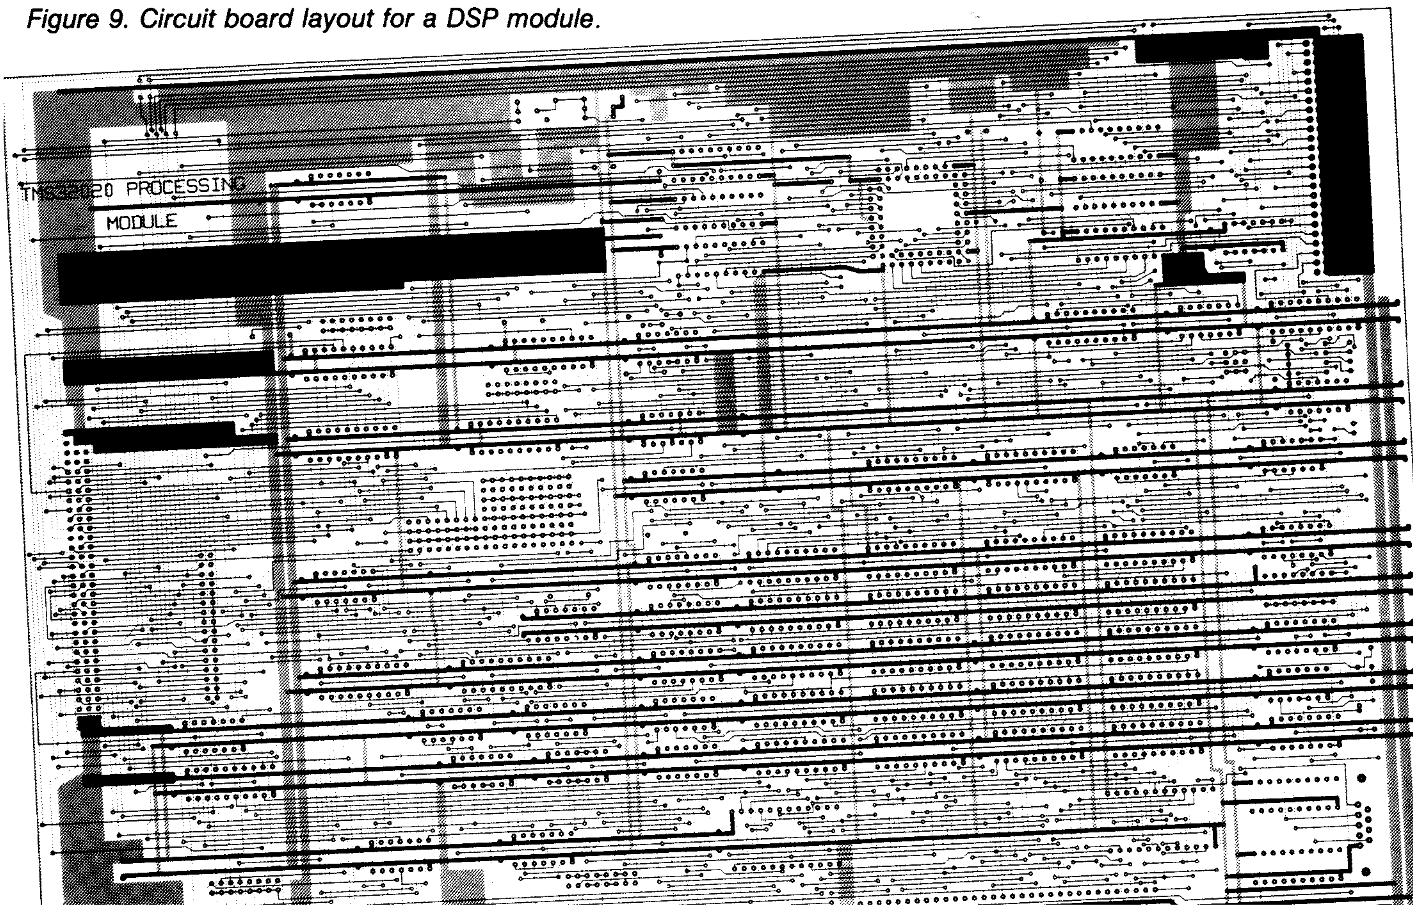

Perspectives

Graeme Gerrard, R.P. Harris and David Hirst

A path in computer music retraced

Graeme Gerrard

Back in 1979, as a composition student at the La Trobe University Music department, I, like most of the other students, was thoroughly embroiled in set theory (*a la* Allen Forte), twelve tone composition, time point theory and the like. After study of Stockhausen's orchestral piece *Gruppen*, I had just completed a tape piece, *Proportions*, that was based on the use of twelve tone manipulation techniques. But I had applied these techniques to durations of events, passages and sections, rather than pitches. This short piece involved numerous measurements of lengths of tape, cutting and splicing and mixing. There was the constant battle against the signal degradation problems inherent in the analogue tape medium. I thought there had to be an easier way; it was just too hard to realise complex pieces with tape. The physicality of the sound as represented by tape offered advantages and (for me) a new way of thinking about sound organisation. But at the same time, the practical side of dealing with sounds on tape was painstaking and time consuming.

At that time Graham Hair, and his student Mark Pollard, were also at La Trobe. They were working with Forte's pitch set data; writing programs to enumerate all the subsets of a given pitch set, I think. For some reason I was very interested in the idea of using a computer to do this sort of thing. It seemed to be appropriate; accurate, relatively error-free, labour saving, repeatable, etc. I knew virtually nothing about computers and, like most people, had only seen them in Sci Fi movies. They seemed to be like a new techno-magic: amazing, but ominous.

I plunged in and taught myself some basic Fortran programming language and started writing little programs immediately; to order lists of numbers, print out "magic squares" of pitch sets etc. Soon I had a few little programs that "composed" music and printed out pages of numbers representing instrumentation, attack times, durations, pitches and dynamics. I would transcribe these into music notation and try to play the results at the piano. I wasn't a bit interested in the computer as a sound synthesiser at the time, but only as a (pre)compositional tool.

That year, my honours year at University, I wrote several music related programs; one that implemented a few of the rules of species counterpoint to generate a melodic line over a given cantus firmus; a program to identify pitch collections in terms of Forte's set types and report on certain properties of the collection; a program to scatter events over a time/frequency matrix using the Poisson formula; a program to enumerate the non-duplicative rotational arrays of pitch sets (a special kind of pitch set transformation used, for example, by Stravinsky in his *Huxley Variations for Orchestra*); and the beginnings of a "serious music composition program" I called COMPOST. This was based on stochastic (random, but generally directed) choices of musical parameters, on a section by section basis. Needless to say, I actually wrote very little music that year, but thousands of lines of code. Computer programming at that time was new, exciting and very satisfying. I was satisfied in solving programming problems, seeing them evolve and "work out", much like a musical composition. An incredible amount of work went into learning programming and writing these programs, but with few exceptions, they were for my own private use and amusement. My view of what I was doing, and why, was very narrow; I was virtually working in isolation. I became interested in sound synthesis, purely as a way of hearing the results of my composing programs without the need to transcribe and play the things on the piano.

Graham Hair had obtained a copy of the MUSIC4BF sound synthesis program, written by Hubert Howe in the mid-1960's, on a previous trip to Princeton University. He had typed it into the computer (a DecSystem10) from a photocopy of a print out, but other commitments prevented him from getting it fully functioning. Brian Parish and myself adapted the program to run on the University's VAX 11/780 computer and Jim Sosnin, lecturer in music technology and sound La Trobe wrote the programs on the Music Department's PDP 11/10 to convert the digital information to sound. The procedure for hearing the result of a synthesis run was a lesson in patience. The digital samples would be calculated and written to a digital tape. We would walk across to the Computer Centre to collect the tape, bring it back up to the Music Department, run a program on the PDP 11/10 computer to convert the samples to analogue form, record that on audio tape and play it back to listen to the result. We were all just learning how to write programs for sound generation and it often occurred that the program generated garbage; noise, silence or something random in between. The turn around time for hearing back music that you synthesised was probably half an hour minimum, for say, a few seconds of sound. It was common to wait overnight to hear a complex synthesis of a few minutes, with the actual calculation time of a sound being of the order of a hundred times the sound's duration. It took a lot of computer time to generate a complete piece, so we often synthesised pieces

in sections or layers and assembled them using multitrack analogue tape. It was an amazingly difficult way of working, but the newness and excitement of it all usually kept us interested. It was a wonder that we ever produced any music at all.

COMPOST was expanding to become a full music composition language. While parameter value selection was still based on various stochastic techniques (e.g. 1/f noise, correlation to specified contours etc.), I had incorporated a method of specifying the organisation of larger syntactic structures. This was based on a hierarchical method, borrowed from Chomsky to describe structure in verbal utterances. It also included a transformational scheme to operate on components of the structure.

In 1982 David Hirst started working on an implementation of the phase vocoder, a powerful system for sound analysis and resynthesis. Brian Parish was working on a graphics-based composition system, and I was working on sound synthesis and generally maintaining and extending the MUSIC4BF program. Eventually, the Music Department acquired its own computer. This was the Charles River 68000 Universe computer and was quite powerful, by the standards of the day. With the help of Jim Sosnin and Brian Parish, I managed to get the MUSIC4BF program to run on the new machine. Unfortunately, the machine was too slow for the computationally intensive task of music synthesis, unless dedicated to a single user at a time.

In 1982 and 1983 we ran summer courses in computer music, with some dozen or so composers each time, including Alistair Riddell, Judy Gunson and Cindy John. These courses ran for about ten days (and nights) and were intended to attract new people to computer music by helping them to realise a small piece each. Judy Gunson and Cindy John were the only ones who were attracted to the medium and went on to write further pieces on their own. But there were several other people who did their own projects as well, Mark Rudolph, for example, implemented his own Linear Predictive Coding (LPC) system from scratch on the Charles River machine. In the time I spent at La Trobe University, I finished about six compositions using the computer music system, including; *Midnight Special* (1981), *Compost* (1981), *Variations for computer* (1982), *Discourse (analogue and digital)*(1982), *Strings of Token Strings* (1984), *Spoken Worlds* (1986), *Harping* (1986), as well as several analogue synthesiser pieces: *Bee Music* (1982), *Streams/Groups* (1983) (both on tape and performed live), *Slang* (1983) (for tape alone, and for tenor saxophone, trumpet bongos and tape), *Solipsis* (1982) (for tape and percussion). Keeping involved with analogue synthesisers was an important way of maintaining an immediate contact with sound, though it lacked the detailed control the computer provided.

In 1986 I started a part-time job in the Faculty of Music at the University of Melbourne, as research assistant with the Computer Music Research Project, initiated by composer Barry Conyngham and Rex Harris from the Department of Computer Science. This project had been operating since the late 1970's and was based around the MUSIC V sound synthesis program, designed at Bell Labs in the early 1960's by Max Mathews. There had been many contributors and participants in the project over the years, including John Chowning, Jean-Claude Risset, Gary Nelson, David Worrall, Amanda Baker and many others (see Rex Harris' article in this issue).

The Melbourne University project had come to the same kind of impasse as La Trobe. It was obvious that software synthesis was of limited use in the University context. It took a long time and a lot of study and work for a composer to become a competent user of these systems; more than most undergraduate or even postgraduate students could afford to devote to learning what is essentially for them, a tool for the realisation of their musical ideas. The problem is mainly technological. General purpose computers are still not fast enough for generating complex sounds in real time. Shared access computing exacerbates the problem. A solution was to turn to the commercially available digital synthesisers, like the Yamaha DX7 and TX816 that, because of the Musical Instrument Digital Interface (MIDI), were able to be controlled by computers.

These synthesisers are capable of executing, in real time, synthesis algorithms far more complex than are usually attempted with software synthesis. The hope was that students would be able to concentrate more time on the event organisation in their music, but still be able to use complex timbres. Unfortunately, these Yamaha devices have a human interface that can only be described as appalling, making it almost universally true that musicians do not design new sounds with these machines, but select from among the numerous available presets. The result has been a superficial approach to timbre in the music that employs them.

The computer composition system in the Faculty of Music consists of access to the University's VAX 11/750, which controls MIDI capable synthesisers via a special MIDI interface. The synthesisers used are Yamaha DX7 and TX816, and an Akai S900 sampler. Using Gary Nelson's MIDIPLAY program, score information is converted to MIDI codes, which are sent to the synthesisers. We have considerably expanded Nelson's original version of MIDIPLAY and included a record mode to capture MIDI performance data. This can be edited and manipulated using a number of programs developed here, including MIDISCORE, ELP (a sort of interactive version of COMPOST), and MPL (Music Programming Library). Having acquired a Mac II computer this year, we are in the process of adapting our software to run on this machine. The MacIntosh was selected because of the utility of its graphic interface for a whole range of teaching, research and music notation applications. Our intention is to develop a music composition work station that makes use of commercially available, special purpose software, MIDI, and a MUSICV/4BF style synthesis language for software synthesis. A short time ago we obtained a Sound Accelerator card from Digidesign. This enables the MacIntosh to generate stereo, CD quality sound, in many cases in real time. Although keeping up with advances in music technology seems to be an expanding spiral, it seems that the minimum requirements of adequate sound quality in a reasonably interactive environment are available. Certainly, there are always going to be newer, faster machines on the horizon (e.g. Steve Job's

NeXT machine has given a lot of attention to sound and music generation in its design), that will provide features approaching one of the goals of music technology, for me, at least. That is, a system that has the responsiveness of a conventional musical instrument, but the flexibility of a general purpose computer that can be adapted to develop in line with my compositional interests.

If the reader has gained the impression that computer music is a hard road, where composers are always rubbing up against the problem of wanting to apply computers to a use for which they were never really designed, that is my perception too. But computer music offers the possibility that entirely new sounds can be created, sounds that nobody has heard before. This opening up of a new sound world, a music where sound itself is the primary focus, is the main motivator. Along with all music, all art, the opportunity exists for changing the world we perceive, seeing it in a new light.

References

Chomsky, Noam.**Aspects of the Theory of Syntax** M.I.T. Press, Cambridge, MA. 1965

Forte, Allen.**The Structure of Atonal Music** Yale University Press, New Haven, Conneticut, 1973

Gerrard, Graeme. **COMPOST: A Language for Computer Music Composition/ Sound Synthesis from Spectral Functions.** M.A. Thesis, La Trobe University, 1986

MIDISCORE Software Documentation. Faculty of Music, University of Melbourne, 1988

ELP — A Musical Event List Processor Software Documentation. Faculty of Music, University of Melbourne, 1987

Howe, Hubert S. Jr. **Electronic Music Synthesis** Norton, N.Y. 1975

Mathews, Max V. et al.**The Technology of Computer Music** M.I.T. Press, Cambridge, MA. 1969

Nelson, Gary.*MPL: A Programming Library for Musical Data Processing* **Creative Computing** 3(2): 76-81. Morristown, USA. 1977

Nelson, Gary. **MPL:/MIDIPLAY Programs** Software Documentation, Computer Music Research Project, University of Melbourne, 1986

Computer music in the University of Melbourne

A brief history

R.P. Harris

Introduction

This paper briefly describes the development of the Computer Music Research Project in the University of Melbourne. The Project is unique in many ways: it demonstrates, for example, successful cooperation between two disparate university departments, one based in the sciences and the other in the arts. Contributions are made by professionals in both fields from their own areas of competence; it is not simply a group of musicians using computers, nor computer scientists playing with music.

No attempt has been made to detail every event, nor catalogue every piece of music or document produced, but rather to indicate in broad outline the major occurrences that have shaped the Project as it is today. In adopting a chronological viewpoint some logical threads may have been disconnected, but perhaps this ensures that events are seen in the context in which they occurred.

Foundation

The Computer Music Research Project was founded in June 1976 as a cooperative venture between the Faculty of Music and the Department of Computer Science. Barry Conyngham, who had recently joined the Music staff, brought with him expertise in the compositional uses of computers, together with a box of punched cards constituting an early version of the Music V synthesis program (Mathews 1969). In seeking collaborators from the computing fraternity he met Jurij Semkiw and the author, the former a skilled hardware designer and the latter a numerical analyst and programmer. With the use of digital to analogue converters borrowed from another project the feasibility of using the Computer Science Department's Perkin-Elmer 8/32 computer for direct digital synthesis of sounds was soon established, and the Project was founded.

The Objects of the Project, as initially formulated, were:

- the direct synthesis of musical sounds
- investigation of the nature of traditional instrumental sounds
- computer assistance in structuring musical works.

To those objectives, which remain essentially unchanged, has been added:

- real-time digital instrument control, with particular reference to developing a system capable of onstage performance.

The early years

The first two or three years of the Project were spent developing the basic computer programs needed for digital synthesis: Music V, a digital sample mixer, spectral analysis procedures, and special programs for a colour graphic terminal.

In order to investigate from a scientific viewpoint simulation of spatially distributed sound, audio equipment such as amplifiers, tape recorder and mixer were configured for four channel working.

While on sabbatical leave in 1977-78 I made several visits to IRCAM, the Institut de Recherche et Coordination Acoustique/Musique in the Centre Georges Pompidou, Paris. The computing section at IRCAM was at that time directed by Jean-Claude Risset. From IRCAM we obtained, among other things, their most recent version of Music V; this formed the basis for a major revision of the program used in Melbourne.

In 1979 there occurred an event that was to provide the Project with a front line resource for musical composition. During July and August Gary Nelson from Oberlin College, Ohio, USA paid us a visit. Nelson, with a doctorate in music composition, is director of Oberlin's music and technology programme and has done much to integrate the use of electronics into music composition and performance. His major contribution to computer music (Nelson, 1977) is the Music Program Library (MPL) system, a collection (still evolving) of over

200 functions embedded in the computer language APL. From the point of view of composers and other musicians, MPLAPL enables the expression and manipulation of structures or gestures directly related to music, rather than requiring the user to write them out laboriously in terms of notes. For a moderate investment in terms of the time needed to become acquainted with this system, the musician has access to a facility that is the equal of anything in the world. While in Melbourne Nelson installed MPL as part of the computer music system, and gave valuable instruction on its structure and use; it is not too much to say that this event revolutionised the approach to using this particular computer music system.

From the Project's inception Music V, under the name 'synth', became somewhat notorious in the Computer Science Department because of the amount of computer time it consumed. To a certain extent, this is unavoidable - generation of millions of samples of a complicated waveform is a computationally intensive process which, like many other computer applications, can utilise to the full whatever computing power is available, no matter how fast. During 1980 the program was completely redesigned and rewritten in the computer language **C**, whose compiler produced much more efficient code than that obtained from Fortran. Music C, as the revised program was named, had the immediate effect of reducing the computation time to about one sixth for the same source score. During this revision considerable care was taken to ensure that there is complete compatibility between the source code accepted by Music C and Music V. Musical scores written for Music V in 1978 can, without alteration, be processed by Music C now (and generate the same sounds).

Finding a place in the world

By this time, the Project had produced some work of lasting consequence. Barry Conyngham made new (and for the first time, digital) realisations of Percy Grainger's *Free Music* (Grainger 1934), two brief but curious pieces intended for a mechanical/electrical musical machine of Grainger's own design. Other music he composed and recorded from the system were *To Be Alone*, and incidental music for the play *No Man's Land*. Research papers on the problems inherent in simulation of the timbre of traditional string instruments (Conyngham & Harris 1982), and on paradigms for computer assisted composition (Conyngham & Harris 1980), were presented at international conferences in Australia and Austria respectively.

A highlight of 1980 was the part the Project contributed to the programme of the International Federation for Information Processing (IFIP) triennial World Computer Congress, held in Melbourne. This contribution took the form of a concert comprising a brief demonstration of the capabilities of the system, followed by the world premier of Barry Conyngham's work *Journeys*. This piece is scored for woodwinds and computer generated tape, and was performed by Peter Clinch playing a range of saxophones and clarinets, with computer sounds reproduced by four loudspeakers arranged in a rather unusual diamond formation.

A significant decision taken about this time involved branching out into the field of real-time computer synthesis. Ever since electronic devices that made musical sounds were first constructed (at least as early as 1920), musicians have sought to use them in the concert hall. Voltage controlled analogue devices have the unfortunate disadvantages of pitch instability and difficulty of control if complex sounds are needed; computers may however generate arbitrarily complicated sounds without stability or control problems. However, Music V and Music C style programs cannot generate sounds in real time, and as a consequence cannot effectively be used on stage. The solution that evolved was to use very fast special purpose hardware for generating the sounds, together with a dedicated processor for control. Examples of these systems found overseas are the $100,000 System Concept Digital Synthesiser (commonly known as the Samson box) controlled by a PDP-10 computer at Stanford University in California, and a much cheaper sequence of machines called 4B, 4C, 4X controlled by LSI-11 and small PDP-11 configurations at IRCAM.

The Melbourne proposal was to use a specially designed and very fast (at that time) bit slice microprocessor for synthesis of sounds, and control it with a PACE microcomputer. The bit slice machine, designated IMI-200, contained several novel design features; had it been delivered in a reasonable time it may well have made a considerable impact both locally and internationally. The PACE however was an unfortunate choice; it came at the end of development of 8-bit microprocessors, and as well as being difficult to use was virtually obsolete.

Gary Nelson paid his second visit to the Project during January and February of 1981, achieving amongst other things the composition of six pieces with the general title *Phrase Structure*. The construction of these pieces in so short a time demonstrated clearly the power of the MPL/Music-C combination in algorithmic composition.

In August Melbourne University was the main venue for an International Music and Technology Conference, held as part of the Victoria State Government's "Music '81" initiative. Both Conyngham and the author were members of the organising committee for this Conference, which had the usual mix of concerts, equipment displays and demonstrations, formal presentations of invited and submitted papers, and less formal discussions of all kinds. Requests are still being received for copies of the Proceedings of this Conference (Harris 1981) which attracted more than 120 delegates including overseas visitors from Britain, Canada, Finland, France, and the USA.

During the period September 1981 to February 1982 I visited the Computer Audio Research Laboratory (CARL) in the Center for Music Experiment of the University of California at San Diego, and the Center for Computer Research in Music and Acoustics (CCRMA) at Stanford University. In the former, I was able to export some techniques developed in Melbourne, and assist with testing the Cmusic system being constructed there at that time by its Director, F. Richard Moore. CCRMA, under its founder and Director John Chowning, is probably the world leader in computer music research.

At Stanford I had access to their latest research results and techniques; indeed, some of this research material has yet to be exploited because of lack of time and resources.

David Worrall, who had joined the Music Faculty in 1980, took an active part in the Project from his arrival. In addition to his teaching duties, he composed *Mixtures and Re-collections* and *Glass Games*, the former for computer alone, and the latter for a small ensemble with computer generated tape.

In September 1982 the Project obtained a new Computer Music Laboratory. After four years of moving from a corner of the tea room to various vacant offices a properly designed listening environment with the loudspeakers at the corners of a square was most welcome.

Also during 1982, after a review of the current state and future prospects for the real time aspect of the Project's work, the Music Faculty provided funds for the first stage of a Unison computer installation. The Unison, an Australian designed and built computer based on the now well-known Motorola MC68000 processor, is large enough to support the Unix operating system, and the Project's machine was among the first to be ordered. While it was being constructed work was started on hardware to connect an IMI-200 to the Unison, and development of PACE software was discontinued.

During 1983 David Worrall continued his compositional activity, producing *With fish scales scattered*, a computer tape piece, and *Silhouettes*, for recorders and computer tape. At about this time he also commenced work developing teaching materials for introducing undergraduates to computer music.

Among Barry Conyngham's compositions were *Voicings* for ensemble and tape, and incidental music for the ABC television series *Into Science*.

Towards a real-time system

After completion of the physical connection between the Unison and the IMI-200, tests soon revealed that the system as it had been designed was too slow to serve its intended purpose i.e. real-time software synthesis of sounds having any degree of complexity; simple voices could be produced, but their musical value was very limited. This problem could probably have been solved to a certain extent by redesigning the software, but external events had by this time rendered the IMI-200 concept obsolete. Commercial digital synthesisers, designed for computer control with sophisicated sound generation modules built into the hardware, were beginning to appear on the market. An industry standard Musical Instrument Digital Interface (MIDI) specification had been formulated, and instruments like the Yamaha DX7 conforming to it were available in the shops. The DX7 as a sound generator is equivalent to at least 50 IMI-200 processors at about the cost of one.

So, early in 1985, the IMI-200 was decommissioned, and a prototype circuit board to enable the Unison to control MIDI devices directly was constructed. Tests with a DX7 borrowed from the Music Faculty produced results so encouraging that an eight module digital synthesiser, a Yamaha TX816, was acquired.

Gary Nelson joined the Project once again in 1985, this time from September until May 1986. Such is the degree of miniaturisation current that he was able to bring with him from the USA his complete personal computer music equipment — computer, synthesisers, amplifiers, etc. As well as developing this system and copying the results to Project resources, he was able to complete some compositions and present a public concert, the last in conjunction with his wife Mary playing the flute.

Expanding MPL to generate MIDI events for synthesisers was a comparatively simple process. As a consequence, much of the work required to create or arrange a piece of computer music is not affected by the ultimate method used for sound production — Music C direct digital synthesis, or MIDI controlled synthesiser. In fact, it is often convenient to audition early versions of a work on the real-time system before using hours of computer time making the definitive realisation with Music C.

However, from this point, compositional activity has had a tendency to use MIDI-controlled devices, and use of Music-C has diminished somewhat. About this time, Barry Conyngham arranged for soprano, ensemble and tape *A life of dreams* from his oratorio *Antipodes* and wrote incidental music for the drama *Spook House*, both using the TX816 synthesiser.

As well as the Unison to MIDI connection, an auxiliary controller that enabled MIDI devices to be attached to a computer terminal port was obtained and programmed, so that TX816 and DX7 synthesisers can be controlled from any of the University's computers. This facility is used only occasionally in the Project, but has found a permanent place as part of the Music Faculty's Electronic Music Studio where it is used for teaching purposes.

The present

With the MIDI connection complete, a limited amount of testing has shown that the system appears to be fully effective as a means of capturing keyboard input and controlling a synthesiser in real time. The strengths and limitations of this configuration will however only be revealed when music requiring its participation is available.

In 1987 and 1988 music composed and presented by Project members has included *Harping* by Graeme Gerrard, who has joined us from La Trobe University, and *Colloquy* and excerpts from *Birdsongs*, both by Amanda Baker.

The future

Although this article is intended to be historical, it is always interesting to guess at what may develop in the next few years. The utility of music performance with real-time computer interaction will continue to be constrained by the speed and versatility of available resources, the bounds of which have not yet been thoroughly explored.

The place of personal computers in a computer music project is still to my mind uncertain. A considerable amount of time (months, if not years) must be invested to develop any system that a professional musician would want to use, and to obtain value from such an investment the system must be attractive to as wide a range of composers as possible, over as long a time as possible. Packages currently available for personal computers do not

appear to satisfy this criterion; most comments I have heard are along the lines that while the composer may be able to make one piece with such a package, they would not want to make more than one. Research, by definition, means pushing at the frontiers of what is known, or is possible, and existing packages do not seem suitable for this purpose.

One advantage of packages of this kind is, of course, that they can be used after a few hours or days of experience, and this is certainly not the case of systems like MPL. There are however sound theoretical reasons why a system that provides a sufficiently general level of facilities must have a certain degree of complexity, and as a consequence require a much longer learning period. For this reason, it seems not to be possible to introduce such a system into the music undergraduate curriculum; time simply is not available.

There are some signs of a swing back to direct synthesis. Although many synthesisers are programmable, it does not seem to be easy to obtain sounds of a kind that some composers appear to want. The ever increasing speed of successive models of computers with larger and larger storage capacities means that direct synthesis techniques may once again become an attractive method of sound generation.

Conclusion

In retrospect, the achievements of the first ten or so years of the Project's existence are quite impressive. For a direct outlay spread over that time of less than $60,000 (the cost of a concert grand piano, or a medium size computer work station), Melbourne University has a computer music facility comparable with the best in the world. While some mistakes and errors of judgement have been made, there have been no major catastrophies and no significant changes in research direction have been required.

A small number of composers have produced a considerable body of music, some of it having achieved recognition world wide. On the technical side, most innovation has been in matters of detail buried inside larger systems, both hardware and software. That the system is not available to a greater number of composers, and that comparatively little fundamental technical research has taken place, are both attributable to the small group of people available for active work at any one time, and the fact that those people have to attend to everything from cleaning the equipment to advanced investigation. This problem is, in fact, universal; only two or three institutions in the world have the resources to employ large teams on computer music research. Indeed, the Project has no full time personnel; all of us associated with it have full teaching and administrative loads, and can only work on computer music in what time remains for research.

The Project is indebted to a great many people who have contributed to its facilities. Most notably, the unique contributions of Jurij Semkiw literally made the Project what it is; without his hardware design and construction skills most of the equipment purchased would have been useless. The Faculty of Music and University of Melbourne Office for Research have been consistent sources of financial support over the years, and the Department of Computer Science has provided floor space, the assistance of its programming and technical staff, and use of its computers.

References

Conyngham, B.E. and Harris, R.P. *Paradigms for Machine Assisted Composition.* **Computer Music Research Project Paper 5** Department of Computer Sciences University of Melbourne 1980.

Conyngham, B.E. and Harris, R.P. *Problems in Computer Synthesis of String Timbre.* In A. Segal (editor), **Proceedings of the Wollongong Co-operative Workshop on the Acoustics of Stringed Musical Instruments,** The University of Wollongong Department of Physics, 1982 pp.1-12.

Grainger, P. 1934 (arr. Barry Conyngham). *Free Music I and II*, Move MS 3027, Move Records, Carlton South, Melbourne 1978.

Harris, R.P. (editor) **Proceedings of the International Music and Technology Conference**. Department of Computer Science, University of Melbourne 1981.

Mathews, M.V. **The Technology of Computer Music** MIT Press, Cambridge, Mass 1969.

Nelson, G.L. *MPL: a Program Library for Musical Data Processing* **Creative Computing** 3 (2) 1977 pp.76-81.

What happened when Australia went digital?

David Hirst

It is now over six years since the first edition of NMA Magazine, in which I related some of my initial experiences in composing for a computer music instrument. At that time I was a part of the first generation of purely digital electronic music workers. I had not come to the electronic medium through analogue synthesis. The following is simply my impression of the last ten years, during which time electronic music has become digitally organised in Australia.

In talking about analogue synthesis recently, Sydney composer Phil Treloar recanted this revelation: "When working with an analogue synthesiser, the sounds are actually running around inside wanting to break out. Composition is the process of releasing them from their bonds."

Phil's comments illustrate the analogue approach to synthesis — subtractive synthesis. Using an information-rich sound source one subtracts the unwanted elements in order to preserve the desired result. A similar approach is used in some soundscapes and musique concrete where environmental sounds are subtracted from one context and placed in another.

When Australian electronic music went digital, sounds had to be constructed from the ground up. The approach to synthesis was essentially additive. The one acception was the Fairlight CMI, but even this instrument had severe restrictions on the size of its storage capacity, and only short sound segments could be digitized.

The additive approach necessitated more rigorous analyses; within sound events (Hirst,1985;

The Australian Network for Art and Technology (ANAT)

to promote, foster and develop the interaction between the arts, sciences and technology

The Australian Network for Art and Technology (ANAT) had its beginnings in 1985 as a 6 month research-based pilot project jointly initiated by the Experimental Art Foundation, Adelaide and the South Australian Ministry of Technology.

It was discovered that art and technology was a significant interdisciplinary growth area, both in Australia and overseas, and required support in order to maximise opportunities for artists working in the field.

In 1987 ANAT collaborated with Artlink magazine on a publication to document new media work involving advanced technology across a range of artforms.

In 1988 the Australia Council devolved its Artists & New Technology Program to ANAT, enabling ANAT's establishment as a separate organisation.

A small seeding fund (the Art Research & Development Fund) was established to directly help artists' projects. This assistance will be continued in 1989.

Artists-in-Residence Programs are being established to provide further access and technical support for artists working in this area.

Other ongoing activities include:

o information collection and dissemination
o national and international networking
o interface with government, industry and education sectors
o development of a data base
o the investigation of new opportunities within the field
o generation & co-ordination of special projects
o promotion of the field through a wide range of media channels

ANAT's quinquennial plan incorporates many of the strands of the earlier pilot project and places a special emphasis on the technological skilling of artists and the facilitation of placements in hi-tech industrial and research environments.

To this end, ANAT is currently co-ordinating a National Summer School in CADCAM (computer-assisted design and manufacture) for artists through the Advanced Technology Education Centre of the Regency College of TAFE in Adelaide. Features of the summer school include individualised tuition based on participants' own creative projects; 24 hour access to the CAD facilities; access to lathe, mill and laser-cutting equipment; and the opportunity to learn computing skills with other artists, craft workers and designers. It is anticipated that ANAT will co-ordinate future short term skilling programs through other technology-based institutions in Australia. Input from artists regarding areas of need, appropriate host organisations, etc, is welcomed.

ANAT's major work project in 1988 was the co-ordination of a small cultural delegation to attend the First International Symposium on Electronic Art (FISEA) held in Utrecht, The Netherlands, in September. Funding was attracted by ANAT for 7 individuals, with another 10 Australians self-funding their participation. The response from the international community to the Australian work presented at FISEA was extremely positive, and links were made to some of the key players in the newmedia scene. ANAT was also represented at the founding of an international society to link existing new media networks and organisations - the Inter Society of Electronic Arts (ISEA). The bid for Australia to host the Third International Symposium was enthusiastically endorsed by ISEA. Initial steps for the symposium, and a major international exhibition to coincide with it, are now being undertaken, and again, input from artists will be sought by ANAT.

ANAT is interested in developing its data base on music/sound/cross disciplinary work & artists in Australia. If you are not on the mailing list and would like to be, please contact: Francesca da Rimini, Executive officer, ANAT PO BOX 21 North Adelaide., 5006. Ph 08 231 9037.

Moorer,1978) and between sound events (Clynes,1982). The computer not only became the means of synthesis, but it was the ideal tool for analysis. There was a preoccupation with computer science, and the Artificial Intelligence community became involved in music analysis (Smoliar, 1980) and wrote counterpoint programmes; analysis of the listening process (Laske,1988,1980; Wessel,1979) and wrote psychoacoustic programmes; analysis of the creative process (Roads,1980) and wrote compositional programmes. The compositions created that were related to this activity were introspective, dealing with inner processes.

The size and cost of computer technology dictated that work had to be undertaken at large educational institutions. Large American institutions like Princeton, M.I.T., Stanford, the University of California at San Diego blazed trails that were followed by the University of Melbourne and LaTrobe University. Their approach was to use a large music synthesis programme (MUSIC V and MUSIC4BF respectively), which was run on a mainframe computer. While these systems were cumbersome, and far from real-time, they were extremely versatile and allowed control of virtually every sonic parameter.

These two systems were both implemented about ten years ago, and the year 1978 was a watershed year for computer music in Australia. Two articles were published that provided the hopes and aspirations of the computer music projects at Melbourne and LaTrobe Universities. Under the title "At Melbourne They Hope to Solve Computer Queries" the New Music Newspaper (1977/78) wrote in an unattributed article:

> The aims of the project are to investigate; (a) the direct synthesis of music by computer, (b) the replacement to some extent of the existing traditional music-making ensembles, and (c) computer assistance of the process of organising and structuring a musical work.

In taking up the third point, Graham Hair's "Prelude to a Project in Computer Assisted Composition" gave us an insight into the directions at LaTrobe. Hair (1977/78) concluded:

> Finally, may I assert that what we are looking at in computer "composing routines" is a question of the possible, not of the necessary or the inevitable. It is surely clear that our age is not an age of a single music practice, but an age of a great number of practices, many of which can easily continue to exist and evolve without the aid of technology, high or low, but solely with those oldest and most well-established of manually-operated machines, musical instruments and voices. To the assertion, whether laudatory or pejorative, that the future of music depends in some way on electronics, computers, synthesisers, technology etc, one can only counter-assert that, on the contrary, it depends as firmly as ever on the perceptual capabilities of human ears and on the conceptual capabilities of minds to which those ears are interfaced.

1978 also marked the release of the recording *Full Spectrum: Australian Digital Music* with works by Grainger/Conyngham, Cary, Burt, Conyngham, and Clynes, using digital devices.

Other institutions utilised dedicated (and expensive) computer music instruments that were becoming available at the time; Martin Wesley-Smith worked with the Fairlight in Sydney, and Tristram Cary used the Synclavier in Adelaide.

Over the next five to six years a number of compositions were released by these institutional leading lights and their students. NMATAPES' first release in 1982 contained five electronic works: *Studies No 1-4 from 'Studies'* by Warren Burt; *Contours, Clowns, and Shadows* by Brian Parish; *Compost(1981)* by Graeme Gerrard; *ONE 1, ONE 1A, TWO 1, TWO 1A* by Alistair Riddell; and my own *Brahmin's Son*.

This was closely followed by NMA's *Computer Music* tape with further works by Riddell, Burt, Gerrard, Parish, and myself, with the addition of *Skidmarks* by Dan Senn, *Electronic Study No. 37(b)* by Martin Wesley-Smith, and *With Fish Scales Scattered* by David Worrall.

A good summary of pieces that were around at the time is provided in the results of the 2MBS-FM "first national competition for radiophonic tape composition" held in 1984/85. Receiving awards were: Peter Mumme of Melbourne for *Veronica Takes a Bath*; Robert Douglas of Sydney for *Homage to Bessemer*; Jon Rose of Sydney for *Colony: Survival in the Right Hemisphere*; Ian Fredericks of Sydney for *Viable Alternative*; Graeme Gerrard of Melbourne for *Strings of Token Strings*; Michael Hannan of Brisbane for *Callisto*; Peter Schaefer of Sydney for *See*; and David Worrall of Melbourne for *With Fish Scales Scattered*.

At this point in time, Australia hadn't registered the full impact of a revolution that began with a meeting held at the National Association of Music Manufacturers conference in the United States in January 1982. The meeting was attended by representatives from a number of commercial musical equipment manufacturers and the result of their meeting was signalled by the first release of a Musical Instrument Digital Interface (MIDI)-standard product by Sequential Circuits in December 1982.

In 1983 Yamaha introduced "a family of programmable digital sound synthesisers based on Chowning's frequency modulation (FM) synthesis". The family's patriarch was the ubiquitous DX7, which was MIDI compatible.

MIDI was designed to enable manufacturers to build compatible equipment. This, combined with the low cost of digital synthesisers,was seen as allowing more people to enter the world of digital synthesis. Musicians would be able to afford their own home studios — a much more egalitarian state. When high-powered microprocessors were added to the market place, it was anticipated that music would be "created all over the planet" in backyard studios, but this doesn't seem to have happened.

Compositional problems have arisen from the standard that was actually created. The restricted nature of MIDI contradicted the advantages of using a computer to make music. While any tuning system is theoretically possible using a computer, MIDI imitated Western popular instrumental music and used the twelve equal subdivisions to the octave. Although this was later rectified by the release of several Yamaha synthesisers with alternate tuning possibilities, MIDI had already killed the computer's creative *raison d'etre*. Commercial designs resulted from a reproductive rather than creative state of mind. This phenomenon was similar to the effect the addition of a keyboard had on analogue synthesis - *Switched-on Bach*.

Many composers did establish home studios, but in some cases this has caused alienation from the "community of inquirers". Working in an institution in the early eighties brought with it a sense of community, a daily sharing of ideas, a stimulus to create. It also served to reinforce certain minority values — a moral support that composers working on their own in the wider community may not experience.

Electronic music composers may also be suffering under the post-modernist purge. This is despite the fact that modernism in music does not necessarily correspond to modernism in the visual arts. Many have argued that modernism in the visual arts was largely created by dealers to sell paintings. If this is the case then their efforts were extremely successful as many dealers and artists retired wealthy. With minor exceptions, this is not the case in art music. Art music created in our time has rarely received popular support. Except perhaps for Cage's notion of the global village, and the work of Nono and Henze, modern architecture and modern painting's ideas of utopia were never directly translated into an acoustical framework — the Futurist "Art of Noises" was quickly suppressed. Yet post-modernism's "classicism-reinterpreted" has become "romanticism-reinterpreted" in the audio domain. Is it because we have already experienced music's neo-classicism? In any event, experimental music is on the ropes.

Restrictions in the output of electronic music composers may also be associated with the problem of reaching an audience. Audiences have never fully accepted an entire concert of "listening to tapes". David Chesworth, Warren Burt and others have attempted to redress the situation with their performances using "live" electronic instruments (less shocking than live electric instruments), however this area is not as well developed in Australia as the work of some performance artists in the United States. Other composers such as Martin Wesley-Smith have written increasingly for acoustic instruments and tape. In France, Pierre Boulez has taken this one step further by incorporating the electronic processing of acoustic instruments in his work *Repons*.

The obvious medium for electronic music is radio, and much more needs to be done to record Australian music and provide access to the airwaves. Film music has proven accessible to the electronic music composer, but much of it has become tepid sonic landscape art.

Institutions need to reassert their role in creating a climate for inquiry and creativity. Perhaps a lack of compatible equipment has seen a fragmentation of effort in this country. While Melbourne and LaTrobe's commitment seems to be wavering between large and small systems, the N.S.W. Conservatorium persists with its Fairlight system, and a number of other institutions, like the University of Tasmania, have assembled MIDI-based systems over a period of time. The introduc-

tion of instruments built around the Macintosh II at a number of centres may encourage the non-duplication of effort and the sharing of software that is developed, whilst retaining the field of all possibilities that the computer can provide. If the larger centres in the United States can form sister relationships between each other and centres such as I.R.C.A.M. in Paris, why should we be so reticent?

The development of live electronic instruments is still an area of great potential and interest. Expo 88 and many other installations, including those of performance artist Leigh Hobba, have highlighted the potential for instruments that interact with the audience. Acoustic environments like Ron Nagorcka's Launceston park project show another side to sonic expression on tape.

Alienation experienced by composers working in the electronic medium could be thwarted by the hosting of periodic gatherings. Whether they are of the grass roots Clifton Hill Community Music type, the academic institutional type, or the industry-based Australian Music Centre type is probably fine detail. What is important is that people with similar interests get together and share ideas, stimulating activity and building minority morale.

References

Clynes, M.(ed) 1982. **Music, Mind, and Brain: The Neuropsychology of Music**. Plenum Publishing, New York.

Hair, G. 1977/78. *'Prelude to a Project in Computer-Assisted Composition'*. **The New Music Newspaper (3)**: 20-21.

Hirst, D.J.G. 1985. **Digital Sound Analysis Using the Short-Time Fourier Transform**. M.A. Thesis, LaTrobe University.

Laske, O.E. 1980. *'Toward an Explicit Cognitive Theory of Musical Listening'*. **Computer Music Journal 4(2)**: 73-83.

Laske, O.E. 1988. *'Introduction to Cognitive Musicology'*. **Computer Music Journal 12(1)**: 43-57.

Moorer, J.A. 1978. *'The Use of the Phase Vocoder in Computer Music Applications'*. **Journal of the Audio Engineering Society 26(1/2)**:42-45.

'At Melbourne Uni they hope to solve computer queries'. 1977/78. **The New Music Newspaper (3)**: 14-15.

Roads, C. 1980. *'Interview with Marvin Minsky'*. **Computer Music Journal 4(3)**: 25-39.

Smoliar, S. 1980. *'A Computer Aid for Schenkerian Analysis.'* **Computer Music Journal 4(2)**:41-59.

Wessel, D.L. 1979. *'Timbre Space as a Musical Control Structure'*. **Computer Music Journal 3(2)**: 45-52.

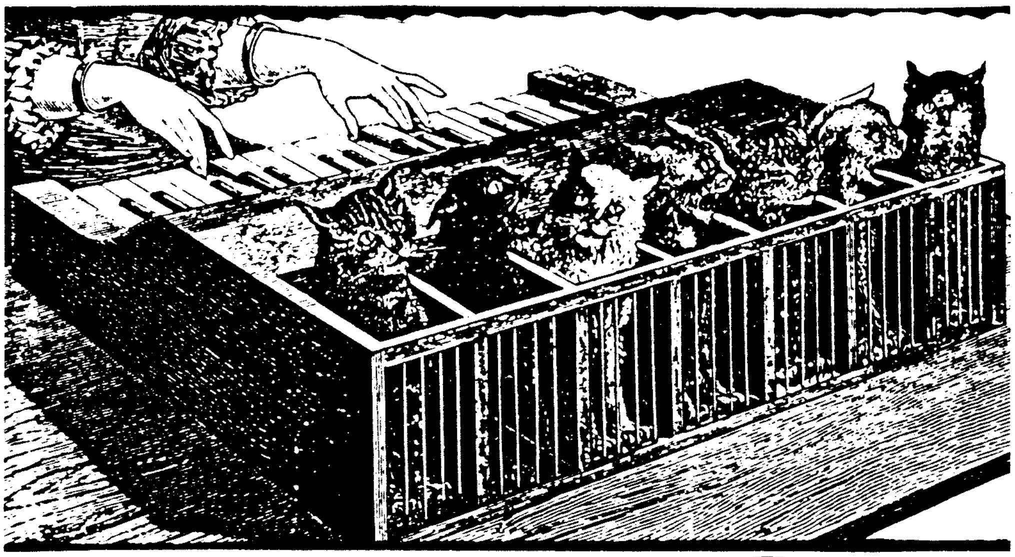

SAmPLe & hOLD